Punch and Judy

A Stress Test for the VSM



Krishan Mathis. Margeret Heath.

Punch and Judy: A Stress Test for the Viable System Model

Krishan Mathis; Margeret Heath

Tautai Press

Draft — 2026-10-26

Copyright © 2026 Krishan Mathis; Margeret Heath
All rights reserved.
Tautai Press
Draft — 2026-10-26

Set in Source Serif 4 with Inter headings.
Typeset with LaTeX (memoir class).

Contents

Prologue: "That's the way to do it" **v**

Introduction **ix**

I Inside: the model we did not read **1**

1 The Mapping We Teach Is Not the Mapping He Wrote **5**

2 System Four Is Not a Pair of Eyes **17**

3 The Law We Quote Is Not the Law Ashby Wrote **23**

4 Channels Without Owners **35**

5 The Ladder Without a Ruler **41**

6 The Brain in the Vat **49**

7 A Machine That Never Became One **55**

8 The Loop That Stayed an Intention **61**

9 Why None of This Ever Cost Us Anything **69**

II Outside: the people who found it without us **77**

10 Six Components, Six Literatures, No Beer **81**

11 The Nesting: Simon, Koestler, Jaques **91**

12 The Variety Logic: Galbraith, Boisot, Aulin, Friston **109**

13 The Autonomy Design: Trist, Herbst, Cherns, Boeke, Endenburg **119**

14 The Evidence: Ostrom **135**

15 The Thirty-Year Experiment We Never Looked At **145**

16 The Objections That Land, and What They Cost **155**

III False Friends and Enemies 171

17 The Market Test 177

18 The Test That Could Not Lose 185

19 A Way of Thinking Has No Failure Condition 197

20 Usefulness Instead of Truth 203

21 The Category Error as Shelter 209

22 The Enemies We Actually Have 215

IV What we lose 221

23 The Narrow Claim 225

24 Six Things Nothing Else Does — and What Each Costs 231

25 The Transducer 237

26 Freedom as a Quantity — The Book He Wrote and We Shelved 243

27 What Might Replace the VSM 255

V What's at Risk in the VSM 261

28 What Would Count Against It? 265

29 The Two Cheapest Decisive Experiments 277

30 The Bill: A Book Plus a Live Register 289

31 Coda: Judy 305

Glossary 311

Acknowledgements 321

Prologue: "That's the way to do it"

"That's the way to do it" is what Mr Punch says after each blow, in the swazzle voice,[1] to an audience that has no vote. It is a verdict announced by the only party keeping score. Any tradition that scores its own results is exposed to the same temptation, and most of all when the tradition is good, because then the self-congratulation is mostly earned. We leave the phrase there and will not press it further.

This book appears in the centenary year of Stafford Beer's birth. Centenaries have a settled shape: collected volumes, retrospectives, panels on legacy, a figure placed carefully in a lineage. That work is worth doing, and much of it is being done well elsewhere. We have chosen a different use for the occasion. Some five decades after the Viable System Model reached the form we now teach,[2] we would rather examine it than commemorate it. Beer built the model by taking things apart. It seems right to mark the year by taking his apart.

We have applied the model to hospitals, ministries, cooperatives, software teams, whole national economies. When an intervention succeeded, we credited the diagnosis. When it failed, we found that the model had been applied badly, or the client lacked commitment, or the recursion had been cut at the wrong level. There is no result we have not been able to accommodate.

That is our charge. It is narrow and, we think, serious: the Viable System Model has never been put at risk. Quantitative tests of it have been published, and Chapter 18 examines the two most substantial; neither was built so that a negative result could cost the model its central claim. It has not been shown to be false. Neither has it been shown to be right, because the community around it, ourselves included, has never stated what observation would make the model give something up.[3] Beer did not ask to be spared. "If mistakes have been made," he wrote, "the models

[1] On the swazzle, the reed a Punch performer places at the back of the mouth to make the voice, and on the history of the show whose catchphrase we borrow: " 'That's the way to do it!' A history of Punch & Judy," Victoria and Albert Museum, last updated 17 April 2024, https://www.vam.ac.uk/articles/thats-the-way-to-do-it-a-history-of-punch-and-judy, accessed 29 September 2026. See also George Speaight, *Punch & Judy: A History* (Studio Vista, 1970).

[2] We date that form from Stafford Beer, *Brain of the Firm* (Allen Lane, 1972), and *The Heart of Enterprise* (John Wiley & Sons, 1979). The second edition of *Brain of the Firm* (John Wiley & Sons, 1981) was prepared in parallel with the latter.

[3] Chapter 9 sets out the criterion we apply here, and Chapter 18 applies it to the published tests of the model.

can be tested — to destruction."[4] He did not say which observation would count, and neither have we. A framework that can absorb any outcome is a language, not a theory. It is useful for describing what we have already decided to see and silent on what we might be wrong about. This thread runs from Chapter 1 to Chapter 30, and it is the only claim in the book we would defend without qualification.

What follows is not an obituary. We are not announcing that the model is dead, and we do not think it is. Nor is it a defense; the community has produced enough of those, and we have written some of them. And it is not a demolition. Part Four, "What we lose," is no courtesy paid at the end. It sets out a narrow claim we believe the model can actually defend: that it is the only available instrument that will take a named system at a named level and say what that system needs in order to remain viable, although we have not operationalized that claim or designed the test that could falsify it. It then names six things we know of nothing else that does: a closed set of functions asserted as necessary, recursion with a criterion, a diagnosis rather than a description, a pathology catalog generated from the structure, one currency across scales and substrates, and the algedonic channel. Chapters 25 and 26 add the transducer as an object of design and autonomy stated as a quantity rather than a value. We state that claim carefully because criticism is only worth reading when it is specific about what it leaves standing. A critique that dismisses everything has told you nothing about which part failed.

Every charge in Parts One and Two ends with a receipt. It names the smallest piece of work that would settle the matter, what that work would cost in people and weeks, and what a negative result would establish, which is the part we have most often skipped. If we claim the neuroanatomical mapping we teach cannot be found in the sources we cite, the receipt is a bounded reading of those sources by someone with no stake in the answer, and a statement of what we will teach differently if the mapping turns out to be there. If we claim the participation loop in our founding case, Project Cybersyn in Chile, stayed an intention and was never built as a channel with standing, the receipt names the archive, the hours, and what it would mean for our account if evidence of such a channel is found. Some of these tests cost a few weeks. The two cheapest decisive ones are in Chapter 29; the most expensive cost rather more, and Chapter 30 prices the whole program. A criticism that names no such test is a complaint, and complaints have never cost this model anything.

We are writing for the people who teach the model, apply it in organizations, defend it in reviews and argue about it at conferences. We are two of them. Between us we have taught the mapping described in Chapter 1 to students who trusted us, quoted the law in Chapter 3 in the form we now believe to be altered, and drawn the channels of Chapter 4 on whiteboards

[4] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 377. The passage continues: "They seem to me to be useful; but they are cheerfully submitted to Popperian criteria of disproof." Chapter 9 takes up what that commitment asks of us.

without noticing what we had left unspecified. The first person plural in this book is meant literally. Nothing here is a charge against a community we stand outside of, and where the record is embarrassing, it is our record.

What we want from the reader is a better version of the disagreement, rather than agreement. If a chapter's charge is wrong, the receipt attached to it should make the error easy to demonstrate. We would rather be corrected in a form we can learn from than granted the benefit of the doubt for another fifty years. The model has survived on the second of those. We think it can survive the first, and we would like to find out.

Introduction

This book began with a failure.

Meg and I (Krishan) were sitting amid the ruins of another book. It was called *Wild Futures* and, at least on paper, it was a perfectly respectable idea. Perhaps too respectable. It had three parts, all of which we thought mattered, but which refused to become one book.

First, there were episodic stories set in possible futures: not as scientific predictions, but as invitations to the imagination. Then came a scholarly engagement with the Viable System Model, the VSM: what has it achieved? What has it overlooked? What questions has it raised that still have not been adequately answered? And finally, there was a political intention. We wanted to talk about agency: about the capacity of people and communities to shape their lives and circumstances. We wanted to ask what systems thinking can contribute to a better life—and whether it is time to read the VSM politically again.[5]

Each of these parts had its place. Together, however, they did not make a book. They stood beside one another like three systems meant to communicate, but speaking different languages. We were looking for a new structure and a name for it. *Punch and Judy* keeps only the second of the three: the scholarly engagement with the VSM.

At some point, Meg said: why not simply *Punch and Judy*?

For her, as an Englishwoman, it was a familiar reference. For me, as a German, it was not. I knew something that seemed distantly related: *Kasperle*, the robber, the crocodile, a gentler world in which the roles are clear and the ending rarely hurts. But at first I heard Meg differently. I heard: *Punching Judy*.

And suddenly another image appeared. Judy as a boxer. Judy taking punches, going down, getting back up, and fighting on. Not because she is invulnerable, but because giving up is not an answer. Perhaps she prevails in the end. Perhaps the refusal to stay down is already a kind of victory.

Both images say something about the history of the VSM.

[5] Earlier political readings include Stafford Beer's own *Designing Freedom* (CBC Learning Systems; John Wiley & Sons, 1974), delivered as the 1973 CBC Massey Lectures; Werner Ulrich, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59, and Beer's reply, "A Reply to Ulrich's 'Critique of Pure Cybernetic Reason'," *Journal of Applied Systems Analysis* 10 (1983): 115–19, which we have not read; and Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011). Chapters 8, 16 and 26 take them up.

In one reading, Punch is the dominant doctrine—the empire of whatever happens to be taken for granted. New ideas appear at the margins; they are absorbed, ignored, appropriated, or simplified until it is difficult to recognise what was new about them in the first place. Systems theory and cybernetics know this experience. They have had great influence, but they have also been misunderstood, reduced, and at times translated into a language that lost their political and human edge.

In the other reading, Judy is the stronger figure. She is not the person who is never attacked. She is the person who does not disappear. She returns, objects, changes shape, and asks us to look more carefully. This is how we want to understand the VSM too: not as a completed intellectual structure, and not as a model to be defended at all costs, but as a living effort to address a question that still concerns us: how do we organise ourselves so that people become—and remain—able to act?

The VSM has achieved a great deal in its history. It has argued that organisations are not merely hierarchies, boxes, or processes, although, as the Prologue says, the argument has never been put to a test it could fail. It has offered a language for thinking about how systems can remain viable under pressure to change. Yet it also has blind spots, open questions, and lines of development that have been explored more fully or supported more convincingly elsewhere. Some issues have received too little attention; other questions have developed in parallel in neighbouring disciplines. Part Two, Chapters 10 to 16, follows those lines into the neighbouring disciplines.

The Beer passages this book cites were located with the help of *Pure VSM*, an unpublished compendium of Beer's statements with their pages that I am assembling. We cite them to Beer's own books, with the printed page, so that each can be checked without it.

This book is therefore neither a commemorative volume nor an indictment. It looks forward. What happens if we turn away from the VSM? What do we lose—and what might we see more clearly? Which of its insights remain indispensable? Which must be revised, supplemented, or left behind?

What made us decide to write it was the centenary of Stafford Beer's birth in 2026. Above all, we want to use that occasion to open a research agenda for the next hundred years. Not because we claim to know all the answers, but because we are convinced that the questions have become urgent again.

Judy, then, stands for more than resistance. She stands for a model that takes blows, changes, and survives precisely because it changes. Over the decades, the VSM has been criticised, simplified, misunderstood, applied, and forgotten. It has moved into new contexts, acquired other vocabularies, and left questions behind that are larger than those it has answers to, so far. The model's viability does not lie in remaining unchanged. It lies in exposing itself to criticism, learning from it, and not losing sight of its central question: how can people and organisations remain able to act

under changing conditions, *while maintaining what they are?* Chapter 6 takes up where the model places that identity, in System Five, and Chapter 7 how the model keeps it fixed.

That means we must not treat the VSM as a monument. A monument preserves a form; a living model has to prove itself under stress, and engage with risk, and perhaps adapt its morphology. It must be able to explain what is different now from the contexts in which it emerged: ecological crisis, new forms of economic power, digital infrastructures, shifting patterns of work and responsibility, and what we observe as a growing experience of powerlessness. It must open itself to insights from other disciplines and allow itself to be corrected where they offer more, or where their evidence is stronger.

We do not want to save the VSM by defending it against every change. We want to take it seriously by thinking about it further. Perhaps that is Judy's real point: she wins not because she stays the same. She wins because she gets up again—different, wiser, and ready for the next round.

That is why this book is called *Punch and Judy*. Not because Punch is the main character, but because Judy gets up again.

Part I

Inside: the model we did not read

Punching Judy

> Judy enters the ring young, with a gift nobody in the gym can name: she sees the whole fight at once. Her footwork came from a diagram handed down in the gym, and no one has since checked it against the original. Her guard drops in places nobody has measured. She has never been scored by a judge from outside her own club.

Each of the nine chapters in Part One isolates one thing we inherited and did not verify. The chapters do three different jobs, and each charge should be weighed by the job it does.

The first job is to set our teaching beside Beer's text. Chapter 1 sets the neuroanatomical mapping we teach beside what Beer actually wrote. Chapter 2 follows System Four as it shrank to environmental scanning, a narrowing that happened in our teaching and not in the source. Chapter 3 looks at the law of requisite variety as we quote it, and at a verb Beer changed between 1966 and 1979. In these three chapters the defect is ours, and Beer's books are the evidence against us.

The second job is to find the places where Beer stated something as a principle, a requirement or an intention, and where neither he nor we made it operational. Chapter 4 turns to the channels between recursion levels. Beer specified a variety balance and an alarm route whose filter the metasystem owns, and left the channels undesigned as institutions, without owners the regulated could name. Chapter 5 presses on the criterion by which recursion levels are identified, which Beer called precise and which, under pressure, resolves to a judgment call. Chapter 6 concerns a model in which people appear as the lowest recursion and as occupants of functions, while what they bring as people was left out of the notation by design. Chapter 7 finds an anatomy of maintenance with no account of becoming: the model describes a viable system in operation and says little about how one comes to exist or changes its kind. Each of these chapters sets out first what Beer wrote, with the page, and then where he and the practice that followed him stopped. The charge is narrower than "Beer never said it," and harder to answer.

The third job is to read the case record. Chapter 8 goes back to the participation loop in our founding case, Project Cybersyn in Chile. Beer's own account documents that loop as an intention and a procedure; so far as we can establish, it was never built as a channel with standing, one that the workers could use against the design.

Chapter 9 asks what the first eight raise together: why none of these has ever cost us anything. Nothing in the model's empirical core has been stated with a condition under which it would count as refuted,[6] and the

[6]The standard is Karl Popper's criterion of falsifiability: Popper and Keuth, *Logik der Forschung*, 11th ed. (Tübingen: Mohr Siebeck, 2005), the edition cited in Chapters 9 and 18. Beer submitted his own models to it: Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley

few conditionals Beer does state are read, when they fail, as evidence of bad diagnosis. So no finding about any of the eight can register as a loss. A defect that cannot produce a cost cannot produce a correction either, and that answer organizes the rest of the book.

These charges are against our own practice. We taught this material, applied it, and did not check it. Each of the nine is stated in a form that can be answered, with a source, a measurement, or a piece of work whose price we name.

& Sons, 1981), 377, where he says of them that "they are cheerfully submitted to Popperian criteria of disproof." Chapter 9 takes up what that commitment asks of us.

Chapter 1

The Mapping We Teach Is Not the Mapping He Wrote

The mapping we repeat

Most of us learned the Viable System Model with a brain attached. The five systems arrived already fitted with anatomy: a system that is the spinal reflex, a system that is the brain stem, a system that looks outward and is therefore, in the shorthand, the eyes. That anatomy is a large part of what makes the model teachable in a morning. It also lends the model an authority it did not earn on its own. Say that System Three is the brain stem and you are offering more than an image. You are implying that a distinction you could otherwise be asked to justify has already been justified by biology.

So the mapping needs an audit, and the audit belongs at the front of this book rather than in an appendix. If the shorthand we teach differs from the mapping Beer set down, there are three possible explanations, and they carry different consequences. Beer may have been wrong about the anatomy. Later teachers may have simplified, substituted or improved his account without recording that they had done so. Or both may have happened, with the second obscuring the first. Argument cannot choose among these. Reading can: the taught text and the written text, put next to each other.

We have not yet done that reading in a documented way. What follows is therefore a charge we cannot yet prove, with the work that would prove or dissolve it. The charge is that the anatomical shorthand circulating in our teaching materials has drifted from the assignments in Beer's own text, and that nobody has examined the drift because no one in the community had reason to look. The evidence that would settle it is unglamorous: dated captures of the sources we actually teach from, quoted exactly, set beside a paginated reading of the founder's text. One of those sources is Jon Walker's guide for co-operatives, written in 1991 and circulated as *The Viable Systems Model: A Guide for Co-operatives and Federations* (version 2.21, 2001) and now as *The Viable Systems Model Guide* (3rd ed., version 3.2, 2020), which gives each of the five systems a part of the body or brain. The Metaphorum VSM

page gives no such mapping; it says only that the model "is inspired by the way the brain co-ordinates the muscles and organs".[1]

The mapping Beer wrote

The reconstruction has to come from Chapters 7 to 14 of the second edition of *Brain of the Firm*, read as text, and not from the diagrams, slide decks or summaries descended from it. That is stricter than it sounds. Much of what the community treats as Beer's account reaches us through intermediaries, and an intermediary that compresses a paragraph into a label has already made an editorial decision on our behalf.

Read directly, the assignments are more specific and less tidy than the shorthand suggests, and they are made across several chapters. Chapter 7 sets out a five-level hierarchy on the vertical axis (pp. 97-99), drawn as five control echelons (Figure 14, p. 99) that, Beer says, "model the somatic nervous system of the body" (p. 98), and the chapters that follow attach the system numbers. System One is the lateral axis "which mediates afferent and efferent information" (p. 98; see also pp. 124, 127). System Two is the spinal cord, its vertebral nodes linked by the sympathetic trunk (pp. 124, 129; see also pp. 108, 176). System Three is the medulla and pons, with the mesencephalon folded in (pp. 98, 140-141), and it handles "the parasympathetic information circuits" (p. 176). Systems One to Three together are Beer's model of autonomic control, "concerned, automatically, with the regulation of internal stability" (p. 76; see also p. 110). System Four is the diencephalon, with the thalami, and the basal ganglia, with the ascending reticular formation as its core (pp. 98, 148). System Five is the cerebral cortex (pp. 98, 144). The receipt at the end of this chapter gives Beer's wording for each. Page numbers here and below are those of the Wiley second edition (Chichester, 1981).

Two features of the written mapping matter before any verdict. Beer names structures, and he groups them by function: the five levels, he wrote, follow "the major functional differences involved" (p. 98). And his assignments are compound: System Four is a set of structures with a stated core, so a teaching label that names any one of them selects from Beer's claim instead of summarizing it. The boundaries between the systems also move from chapter to chapter. On p. 132 the sympathetic loop appears in Beer's account of "The autonomic system, number Three"; on p. 176 the same circuits are "the sympathetic (System Two) circuits". The hypothalamus is "the highest element in System Three, or the lowest element in System Four" (p. 141), and the ascending reticular formation, "the core of System

[1] Jon Walker, *The Viable Systems Model: A Guide for Co-operatives and Federations*, version 2.21 (last modified 11 November 2001; original version October 1991), https://metaphorum.org/wp-content/uploads/2022/11/vsm-coop-walker.pdf; Jon Walker, *The Viable Systems Model Guide*, 3rd ed., version 3.2 (2020), https://vsmg.lrc.org.uk/; Metaphorum, "Viable System Model," https://metaphorum.org/staffords-work/viable-system-model. All accessed 29 September 2026.

Four" (p. 148), runs "out through the highest structure of System Three — the mesencephalon" (p. 141). Whether the selection distorts him is exactly what the audit is for.

Beer also said how much the count itself should carry. Introducing the five levels in Chapter 7, he described "what seems to be a five-level hierarchy of systems" and explained the hedge: "I say 'seems to be' because the choice of five (rather than some other number) is somewhat arbitrary" (p. 98). On his own account, the anatomy did not settle how many systems there are.

He also said what kind of claim the mapping was, and it was more than an illustration: "Models are more than analogies: they are meant to disclose the key structure of the system under study" (p. 75), and "There can be no interest in analogy at this point" (p. 87). Yet the body was one choice among several. Beer could have used "the amoeba, or a whole animal species" as the model, and "the results are the same" (p. 75). In *Decision and Control* he called the brain model "a homomorphic mapping of the brain which is by definition a simplification," and "a simplification of which the limitations are understood," and he held that a model which leaves out relevant elements "will lose in utility, but it cannot lose in validity."[2] The audit therefore has to say which claim it tests. On Beer's own account an anatomy that fails costs the model usefulness, not validity.

Whether Beer's neuroanatomy survives current science is a later question, and a different one. The prior question, the one the community has skipped, is what he actually claimed, since we cannot evaluate a mapping we have not first stated correctly. Chapter 2 takes up what happened to System Four, where the compression has gone furthest. Chapter 3 shows the same pattern at work on a borrowed law instead of a borrowed anatomy.

The shelf behind the diagram

Beer's bibliography is partly system-tagged, which makes it unusually useful evidence: for some entries it tells us which reading he associated with which part of the model. Under the heading "Specific Neurological Systems" in the Select Bibliography (pp. 405-407), the entries by Mitchell and by Eccles, Ito and Szentágothai are listed "On Autonomics (Systems One-Two-Three)", Brodal "On 'arousal' mechanisms (System Four)", and Burns and Sholl "On the Cortex (System Five)". General entries such as Netter and Magoun carry no tag, and no source is tagged to System Two or System Three alone. A short account of what each supplied to the mapping

[2] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966; Stafford Beer Classic Library reprint, 1994), 136 (homomorphic mapping) and 113 (utility and validity), pages as in the 1994 reprint. At 136 Beer adds that "the really pejorative term here would be over-simplification; but that stage is not reached so long as a certain structure is preserved in the homomorph, and so long as we know exactly what that structure is." Pages read 29 September 2026 in the PDF.

would let us reconstruct the materials that were actually on the desk when the mapping was made.

Three things are easy to collapse into one here, and they need to be kept apart: the anatomy Beer explicitly described, the scientific literature he explicitly cited, and the later stories readers have projected backward onto him. The first is a matter of quotation, the second a matter of the bibliography. The third is a matter of our own habits, and we are least equipped to notice it, because a story we have been telling for years feels like something we read.

The bibliography does not establish that Beer read his sources correctly, or that he used them as their authors intended. What it does establish is more modest and more useful: the range of positions available to him, and the vocabulary he had reason to adopt.

The Beer we should not invent

A suspicion comes up often enough in our own conversations: that the model rests on a crude layered-brain story or on a simple reflex arc, and that this is why the anatomy now looks dated. We tested it before repeating it, because criticism aimed at a position the author did not hold is worse than none at all. It discredits the critique and protects the original.

The reflex-arc half fails in Beer's own words. Chapter 8 prints the arc as Figure 15, a "Familiar simplified account of the reflex arc" (p. 104), and then sets it aside: "Even so, this reflex arc explanation will not do, because it is too simplified" (p. 105). The layered-brain half fails as far as a search of the text can take it. We searched the same edition for MacLean, "triune," "limbic," Sherrington, Pavlov and Moruzzi, and found none of them. A negative search is only as good as its record, so the record is part of the result.

- **Edition searched:** *Brain of the Firm*, 2nd ed., John Wiley & Sons, Chichester, 1981 (calibre copy, a scan of the printed book)
- **Search surface:** both: the full text, including glossary, Select Bibliography and index, searched through the scan's OCR text layer
- **Search date:** 2026-09-23 (OCR text layer); 2026-09-29 (printed index, read in the page images)
- **Search terms and variants:** case-insensitive substrings MacLean (also "Mac Lean"), "triune," "limbic," Sherrington, Pavlov (covering "Pavlovian") and Moruzzi, each also as a truncated stem to allow for OCR errors
- **Result:** MacLean: 0 hits. "Triune": 0 hits. "Limbic": 0 hits. Sherrington: 0 hits. Pavlov: 0 hits. Moruzzi: 0 hits. (Magoun, Moruzzi's co-author of 1949, appears once, as the Select Bibliography entry for *The Waking Brain*, p. 406. Moruzzi's absence says little: Beer cites Magoun's 1958 synthesis and builds on the ascending reticular formation, "the core of System Four" (p. 148).) The printed index

(pp. 411-417) was also read page by page in the scan, to catch what the OCR layer might miss: none of the six terms has an entry, and neither has Magoun

The finding is precisely this and no more: these names and formulations do not appear in the stated search surface. Absence in a text is not absence from a mind. Beer worked in a period when several of these ideas were ambient, and he could have absorbed a framework without naming its author. What the search can do is stop us attributing to him a specific doctrine he did not invoke. That aim is narrow, and it is the aim of the whole book: to criticize the mapping he wrote, and leave alone the more convenient mapping assembled for him afterward.

What changed in the brain

Only now does the scientific question become answerable, and it has to be answered element by element. "Outdated" has to be earned for each structure; it cannot be a mood applied to a text of its age. For every anatomical claim we ask whether the named structure is still described substantially as Beer described it; whether its function has been revised, distributed or subdivided; whether current work rejects the assignment outright; whether the claim has survived but become incomplete because later work added mechanisms Beer could not have known; and whether the revision matters for the organizational function assigned to that VSM system.

The last question keeps the audit honest, and it requires holding two verdicts apart: the standing of the neuroscience, and the consequence for the model. A biological analogy can fail while the organizational distinction it illustrated remains sound. Then we have lost an illustration and kept the argument. An anatomical claim can also survive intact without validating the organizational use made of it, since the anatomy never licensed that inference in the first place. Collapsing the two verdicts is how a metaphor acquires the standing of a finding, and it is how we have been reading this mapping. Chapter 9 asks why the collapse has cost us so little for so long.

Receipt: an audit we can publish in months

What this chapter delivers is narrower than another general verdict on the use of metaphor in management cybernetics: a bounded, reproducible audit that a small group could complete in months and that anyone could then check or overturn. It covers all five systems. The current-neuroscience column, and the status codes that depend on it, are to be commissioned from someone competent to read that literature. The commissioned reader starts from one named outside text, *Principles of Neural Science* (6th ed., 2021), and may confirm or replace it.[3]

[3] Eric R. Kandel, John D. Koester, Sarah H. Mack and Steven A. Siegelbaum, eds., *Principles of Neural Science*, 6th ed. (New York: McGraw Hill, 2021).

VSM system	Anatomical element in Beer	Beer quotation and page	Source Beer tagged to the system	Mapping in current teaching source	Current-neuroscience finding	Status	Consequence for the VSM claim	Evidence still needed
System One	Lateral axis (afferent and efferent nerves) at the spinal vertebral level	"the lateral axis which mediates afferent and efferent information" (p. 98); "This, then, is what the lateral axis, System One, has to do" (p. 127)	Mitchell 1953; Eccles, Ito and Szentágothai 1967 (joint tag "Systems One-Two-Three")	[collect]	[research]	refuted / revised / unaffected / superseded-with-addition	[state narrowly]	[list]

VSM system	Anatomical element in Beer	Beer quotation and page	Source Beer tagged to the system	Mapping in current teaching source	Current-neuroscience finding	Status	Consequence for the VSM claim	Evidence still needed
System Two	Spinal cord; sympathetic trunk linking its vertebral nodes	"the analogue of the spinal cord" (p. 124); "the analogue of the sympathetic trunk, which links together the vertebral nodes of System Two" (p. 129)	Mitchell 1953; Eccles, Ito and Szentágothai 1967 (joint tag "Systems One-Two-Three")	[collect]	[research]	refuted / revised / unaffected / superseded-with-addition	[state narrowly]	[list]

VSM system	Anatomical element in Beer	Beer quotation and page	Source Beer tagged to the system	Mapping in current teaching source	Current-neuroscience finding	Status	Consequence for the VSM claim	Evidence still needed
System Three	Medulla and pons, with the mesencephalon: the parasympathetic (autonomic) circuits	Echelon statement (Control Echelon III): "I see no reason why this should not be included with the medulla and pons" (p. 98). System statements: "the sleep centres in the pons and medulla" (p. 140); "the highest structure	Mitchell 1953; Eccles, Ito and Szentágothai 1967 (joint tag "Systems One-Two-Three")	[collect]	[research]	refuted / revised / unaffected / superseded-with-addition	[state narrowly]	[list]

VSM system	Anatomical element in Beer	Beer quotation and page	Source Beer tagged to the system	Mapping in current teaching source	Current-neuroscience finding	Status	Consequence for the VSM claim	Evidence still needed
System Four	Diencephalon (with thalami) and basal ganglia; ascending reticular formation as core	Echelon statement (Control Echelon IV): “the diencephalon and basal ganglia (our fourth level)” (p. 98). System statements: “the great linking mechanism between volitional and autonomic control. It is the biggest	Brodal 1957; Kilmer, Blum and McCulloch 1969 (“‘arousal’ mechanisms (System Four)”)	[collect]	[research]	refuted / revised / unaffected / superseded-with-addition	[state narrowly]	[list]

VSM system	Anatomical element in Beer	Beer quotation and page	Source Beer tagged to the system	Mapping in current teaching source	Current-neuroscience finding	Status	Consequence for the VSM claim	Evidence still needed
System Five	Cerebral cortex	"the cerebral cortex itself" (p. 98); "their highest — that is, System Five — representation in distinct and different locations of the cerebral cortex" (p. 144)	Burns 1958; Sholl 1956 ("On the Cortex (System Five)")	[collect]	[research]	refuted / revised / unaffected / superseded-with-addition	[state narrowly]	[list]

The four status codes are defined before use, and each must be defensible on the row's own evidence:

- **Refuted:** current evidence contradicts the anatomical or functional assignment.
- **Revised:** the assignment retains a defensible core, but its formulation must change.
- **Unaffected:** the relevant claim remains serviceable at the level of precision Beer used.
- **Superseded-with-addition:** the claim is not simply false, but later findings add mechanisms or structures that materially change the comparison.

The method used here is the method of the book. Take an inherited claim, recover the founder's actual text, consult the current science, and specify the smallest piece of work that would settle the matter. Part Two applies it where neighboring literatures have moved on without us; Part Four asks what we lose by the delay; Chapter 28 states what would count against the argument made throughout.

We have taught a compressed mapping, inherited its authority and rarely checked the text beneath it. That is a description of ourselves. The correction begins with an exact account of what Beer wrote, continues with an element-by-element scientific audit, and ends with either a version of the mapping we can defend in front of a neuroscientist or a metaphor we should stop using.

Chapter 2

System Four Is Not a Pair of Eyes

The pair of eyes we made

When we teach the Viable System Model, System 4↗ usually arrives as a pair of eyes. It looks outward, scans the environment for threats and opportunities, and reports what it finds to the rest of the organization. The picture is compact and can be drawn in a few seconds. It also maps without friction onto a department that already exists under the name of strategy, intelligence or market research. That is why it travels.

We think this is a reduction, and more than a convenient shorthand. The case for that has to rest on evidence: several current teaching formulations set side by side, quoted precisely, with the scope and date of the search recorded so that a reader can check whether we sampled fairly. We assembled a first set on 24 September 2026: 25 independent web texts, four practitioner textbooks and eight pages of our own training material. Thirteen of the 25 web texts describe System 4 chiefly as scanning or looking outward,[1] and only three say that it holds a model of the organization itself. Our own training pages follow the scanning pattern on all eight pages checked. It is a convenience sample from one search engine, so the claim describes what we have taught and what we found taught; it does not measure the field.

We write this in the first person intentionally. We did not simply inherit the reduction from a careless secondary literature; as a community we reproduced it. Eyes are easy to draw, easy to teach and easy to assign to an existing unit, and Chapter 9, *Why None of This Ever Cost Us Anything*, describes the conditions under which a simplification like this persists without correction. Its cost is specific. When sensing is presented as the whole of the function, modeling reads as an optional activity, something a

[1] Beer's name for the function is the "outside-and-then" (Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 111) and the "development directorate" (*Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 181; *The Heart of Enterprise* (John Wiley & Sons, 1979), 235–36) — read. "Environmental scanning" is the term in strategic management (Francis J. Aguilar, *Scanning the Business Environment* (New York: Macmillan, 1967); H. I. Ansoff, "Strategic Issue Management," *Strategic Management Journal* 1, no. 2 (1980): 131–48) and "horizon scanning" the named practice in futures studies (E. Amanatidou et al., "On Concepts and Methods in Horizon Scanning," *Science and Public Policy* 39, no. 2 (2012): 208–21; T. Könnölä et al., "Facing the Future: Scanning, Synthesizing and Sense-Making in Horizon Scanning," *Science and Public Policy* 39, no. 2 (2012): 222–31) — not read. One text in our sample uses the phrase itself: Dave Hulbert, "Wardley Leadership Strategies," 12 October 2025.

different department performs afterward, and the question of what System 4 is for drops out of the teaching.[2] Chapter 1, *The Mapping We Teach Is Not the Mapping He Wrote*, sets out the general pattern. What follows traces one instance of it in detail.

What Beer intended System Four to do

Beer's own pages give System 4 a larger job than watching. In *Brain of the Firm* he writes that System Four "must contain a model — some model — of the corporation", and in *The Heart of Enterprise* that it "must contain a model of the viable system of which it is the System Four". That model is where the organization keeps its account of the environment and of its own position in it. A model of this kind does more than record what is currently the case. It forms expectations, checks incoming signals against them, and revises its account fast enough for the revised version to be usable in action; Beer calls this foresight by "rapid-time simulation". Scanning is one input to that work and does not describe the whole of it.[3]

The same passages change how the loop between System 3 and System 4 can be read. In the usual reading the loop is a hand-over: System 4 reports what it has seen outside, and System 3, with System 5 arbitrating, turns the report into instruction for the operations. Beer describes something less one-directional. In *Brain of the Firm* he speaks of a "homeostatic loop between sensation and action" and of a feed-forward loop whose predictions "are compared with real-life outcomes"; in *The Heart of Enterprise* the intended impact of an outward act "is fed back" to System Four's model of itself. If those passages carry the weight we put on them, what System

[2]The claim is about what the teaching names, not about what the system does. In the system the two are not separable: System Four is "concerned to manage the outside-and-then, but to provide self-awareness for the System-in-focus" (Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 115), and "it is just this infinite regression of self-images that seems to hold the key to the characteristic self-awareness of viable systems" (ibid., 116) — read. What drops out of the teaching is the name for the modeling, not the modeling.

[3]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 182–83: "the homeostatic loop between sensation and action"; "System Four must contain a model — some model — of the corporation"; a corporate plan "must continuously abort"; and 189, a "feed-forward" loop whose predictions "are compared with real-life outcomes." Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 235: "System Four must contain a model of the viable system of which it is the System Four"; and 240, where the "intended impact" of an outward act is fed back to "the System Four model of itself." Two terms here belong to Beer rather than to current usage. "Homeostatic" is apt for the Three–Four homeostat, in which each system controls the other (see the glossary), but the relation Beer calls a homeostatic loop between sensation and action holds no set point and is not homeostasis in the physiological sense; the modern account of the same relation describes gain control and selective release, not a maintained equilibrium (see the action-selection section below). And the plan that "must continuously abort" is not a method from operational research: Beer gives his reason in the same sentence — "unless that plan is continuously aborting, how shall sanity (the match between the world and the firm's notion of it) be preserved?" — that is, the plan is held provisionally so that outcomes can overturn it, which is a feedback requirement and not an optimisation.

4 sends into the homeostat is a prediction about what the operations will show next, and what System 3 sends back is the check. An instruction issued on the strength of that prediction is a wager, framed so that the next round of operational reporting can settle it. Prediction and report then sit on the same loop rather than at opposite ends of a hierarchy.[4]

We state this conditionally because the step from Beer's passages to "the homeostat as a checkable wager" is ours, and we will not claim more for him than the passages warrant. The phrase "prediction channel" appears in none of the eight books we searched. One further temptation needs resisting. The predictive-processing vocabulary of current neuroscience fits this reading closely enough to be seductive. Where we use it, the comparison is ours and says nothing about what Beer wrote or intended.

The anatomy carrying the argument

Beer's account of regulation drew on the neurophysiology available to him, and part of that neurophysiology is load-bearing rather than illustrative. The item at issue is a reticular-formation model associated with Kilmer, Blum and McCulloch and dated to 1969. We name it because the research brief names it, and because Beer names it himself: the bibliography of the second edition of *Brain of the Firm* lists, under arousal mechanisms (System Four), W. L. Kilmer, J. Blum and W. S. McCulloch, *The Reticular Formation* (Air Force Office of Scientific Research, Arlington, Virginia, 1969). What that report actually claims remains to be verified. Until it is, we are naming a target for checking, not citing a source.

We decline the easy argument, that old anatomy invalidates the organizational logic built on it. Chapter 28, *What Would Count Against It?*, sets out the distinction we rely on here. An analogy can be decorative, and then replacing it costs nothing. Or it can be load-bearing, and then something in the organizational argument moves when the substrate moves. Which of these applies to the reticular-formation model has to be established passage by passage, and not assumed in either direction.

Set against that older substrate is the account described in the research brief: an arousal architecture running from parabrachial nuclei to the basal forebrain, in which monoamines act as gain controls rather than as a single energizing command. The regulatory logic differs, and the difference matters before any organizational consequence is drawn. On the older

[4]The phrase "prediction channel" occurs in none of the eight Beer books we searched (see the note on the search below). The reading of the System 3–System 4 homeostat as a checkable wager is therefore ours, supported by Beer, *Brain of the Firm*, 182–83 and 189, and Beer, *The Heart of Enterprise*, 240. The sentence that described "the channel between System 4 and the command channel" was dropped on 12 October 2026, in the rewrite of this section around the System 3–System 4 homeostat: Beer routes System Four's indications into System Five, not into the command axis (Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 181: "they too are switched into System Five"), and the adaptation link he names is the Three–Four homeostat (Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 120: "The Three-Four homeostat is the organ of ADAPTATION for the enterprise").

picture, a command signal switches a system into a state. Gain control adjusts how strongly many distributed processes respond to their own inputs, so that what changes globally is sensitivity rather than instruction. If the organizational inferences we teach depend on the first picture, they need re-deriving under the second.

Action selection: the function the diagram omits

A system that senses and models still faces a problem neither sensing nor modeling solves. Most of the actions available to it are mutually incompatible, and something has to decide which one is expressed. This is action selection, and the five-system diagram does not name it.

The comparative literature describes the vertebrate solution in terms we are borrowing, so we need to state them accurately. The output nuclei of the basal ganglia are tonically inhibitory: at rest, they suppress. The default is therefore suppression of the available actions rather than quiescence, and selection occurs when inhibition is released from one candidate while the others stay held back. Grillner and Stephenson-Jones are the starting point we have been given for the comparative lineage. The claims about tonic inhibition, release from inhibition and homology each need checking against named sources, and should not travel as a package.

Pictured this way, the central regulatory problem changes shape. The question becomes how a field of possible actions is held back, compared and selectively released, where before it was how a command center issues an action. That architecture differs from a pair of eyes wired to a vertical command line, and it is not obvious where the VSM puts it.

The evolutionary depth of this arrangement matters only if it can be sourced. A claim in circulation says the logic has been continuous for approximately 560 million years. If that survives verification, action selection is an ancient organizing problem that any viable system has had to solve, and no late refinement of cognition. We record it here as a claim under verification. If it does not survive, the number comes out; we will not soften it with a hedge.

The model without a forward model

Our textual search of the defined Beer corpus returned a negative result: no function distinct from System 4 appears to perform forward modeling, in the sense of predicting the sensory consequences of the system's own action. We searched eight Beer books on 24 September 2026 in the OCR text of our calibre copies.[5] "Forward model" and "prediction channel" occur

[5]The search covered the OCR text layer of our calibre copies of *Brain of the Firm*, 2nd ed. (1981); *The Heart of Enterprise* (1979); *Diagnosing the System for Organizations* (1985); *Decision and Control* and *Platform for Change* (1994 reprints); *Designing Freedom* (unpaginated transcript); *Think Before You Think* (2009); and *Beyond Dispute* (1994). Terms: "forward model", "prediction channel", "feed-forward", "predict", "anticipat", "foresight" and "simulat". "Forward model"

in none of them. The function, however, turns up inside System 4 under other names, as foresight by "rapid-time simulation" and as a feed-forward model whose predictions are compared with outcomes. The result stays provisional: OCR can hide hits, and the 1972 edition of *Brain of the Firm* was not searched.

Taken at face value, the finding leaves two options, and both cost something. Either forward modeling already belongs inside System 4, and then teaching System 4 as environmental scanning has quietly removed one of its central operations from what we pass on. Or it does not belong there, and then the five-system architecture lacks a separately specified function that a system needs in order to anticipate the consequences of its own action. We are in no position to choose between these by analogy with neural architecture, and we do not propose to.

Change in the world, not change of kind

System 4's anticipation has an outer limit. It equips a system to model changes in its environment while it remains the kind of system whose viability the model exists to preserve. By itself, it does not account for a system becoming a different kind of system.

That boundary is all we claim. The morphogenesis argument belongs to Chapter 7, *A Machine That Never Became One*, and we do not restate it. We mark where System 4's competence ends only so that the dependency analysis below is not asked to carry a question it was never framed to answer.

The receipt: what falls with the substrate

Each claim below is listed on its own and put as an open question about what survives if the 1969 substrate is replaced by the modern account. We do not answer the questions here. Answering them is the work this chapter asks for.

1. **System 4 must sense the environment.** Does this follow from the requirements of viable regulation independently of the 1969 anatomy, or does the argument as we teach it route through that anatomy?
2. **System 4 must maintain a model.** Is this secured by Beer's own argument, by the good-regulator tradition and the internal-model principle examined in Chapter 3, *The Law We Quote Is Not the Law Ashby Wrote*, or by none of these without further conditions being supplied?

---

and "prediction channel" have no hits. The function appears in Beer, *Brain of the Firm*, 183 (foresight by "rapid-time simulation"), 189 and 196 (fast-time simulation), and Beer, *Heart of Enterprise*, 229 and 240. For the term as defined in motor control (abstract only): R. C. Miall and D. M. Wolpert, "Forward models for physiological motor control," *Neural Networks* 9 (1996): 1265–79.

3. **System 4 must anticipate.** What kind of prediction is actually required, and what evidence establishes anticipation as a distinct organizational function rather than a description of competent modeling?
4. **The System 3–System 4 homeostat carries predictions, not only reports.** Does this reading survive replacement of the reticular-formation model, and if so, on what warrant?
5. **The five-system partition is complete.** Where does action selection reside, and does locating it inside System 4 preserve the partition or turn System 4 into an unconstrained remainder that absorbs whatever the other four do not cover?
6. **Neural organization licenses organizational design.** For each claim that survives, is the warrant formal, empirical, analogical, or merely suggestive?

Replacing an obsolete substrate need not abolish System 4. The likelier outcome is a more exact account of modeling, prediction and action selection than the pair of eyes ever gave us, with a clearer statement of which parts of that account rest on argument and which on borrowed anatomy. Part Five takes up what such a reconstruction would require of us, and the Coda, *Judy*, asks what we owe the model once the work is done.

Chapter 3

The Law We Quote Is Not the Law Ashby Wrote

Ask whether anything in our tradition is proved, and the answer comes back quickly: Ashby's law. Only variety can absorb variety. We have quoted it ourselves, many times, with Ashby's name under it.

Ashby did not write that sentence.

In chapter eleven of *An Introduction to Cybernetics*, 1956, he wrote: "only variety in R can force down the variety due to D; variety can destroy variety."[1] The version we quote is Beer's.[2] It is the better sentence, good enough to travel on its own, and it left behind the conditions that made Ashby's result exact.

The good regulator theorem, which we lean on almost as hard, has a different problem. It does not prove what we cite it for, and its authors said so in the paper itself, in print, in 1970.

The law holds, and the sections below say why it holds and how far it reaches. What needs correcting is how we read it and what we have attached to it. The correction asks for no experiment, no fieldwork and no revision of the model. It asks for different footnotes.

[1]W. Ross Ashby, *An Introduction to Cybernetics* (Chapman & Hall, 1956), 207. The sentence is in section 11/7, where the law is stated in logarithmic form and Ashby introduces it as the way to put that law "more picturesquely"; it does not close the section — one sentence follows it before 11/8. Section 11/6 has already proved the non-logarithmic version and closes on a near-identical sentence: "Only variety in R's moves can force down the variety in the outcomes."

[2]"Only variety can absorb variety" is Beer's compression rather than Ashby's sentence. See Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), in the chapters on the Chilean project: "Only variety can absorb variety: Ashby's Law can be met either by expanding regulatory variety to absorb evolutionary variety, or by curbing evolutionary learning until variety in the economy matches the regulatory variety disposed by the only regulator... that the ideology of the status quo is prepared to acknowledge." The sentence is at page 308, in the section "Externalities"; page checked 22 September 2026 against a scan of the Wiley second edition. Beer had quoted Ashby's own verb in 1966: the law "may be expressed (in Ashby's own words) 'picturesquely: only variety can destroy variety'" (Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966; Stafford Beer Classic Library reprint, 1994), 279, page as in the 1994 reprint). Beer's quotation has "only"; the reprint of Ashby's page 207 cited in the note above does not. By 1979 the verb is Beer's: in *The Heart of Enterprise* (John Wiley & Sons, 1979) he introduces "VARIETY ABSORBS VARIETY" as "the first stage of our perception about variety regulation" (86), writes of "the natural law that variety absorbs variety" (88), and twice sets the *absorb* form under Ashby's name: "These diagrams are embodiments of Ashby's Law. They say: only variety can absorb variety" (286); "Ashby's Law: ONLY VARIETY CAN ABSORB VARIETY" (497). Pages read 29 September 2026 in the PDFs.

None of this is an argument against Ashby, against Beer, or against the Viable System Model. It is an argument about our reading habits.

Two sentences and one verb

Put the two verbs side by side. *Destroy* describes what a regulator does to the variety that reaches it. *Absorb* describes what the regulator must be able to hold. With Beer's verb the law becomes an instruction to a designer: match the varieties, by amplifying the regulator or attenuating what it regulates.[3] Ashby's verb gives no instruction. It states a bound.

A compression is not a falsification. In one sense Beer's version is faithful to Ashby. Anyone who accepts the theorem should accept the design advice. What the compression drops is the apparatus: the stipulation the proof rests on, the constant that appears when the stipulation is relaxed, and Ashby's explicit statement that the result holds *a priori* and so says nothing about any particular organization. Quote the compression and give the 1956 book as its source, and two things happen at once. Beer speaks under Ashby's name, and a design instruction is read as a bound.

This is ordinary usage. Martin Pfiffner's *The Neurology of Business*, 2022, a book on implementing the Viable System Model, sets "Only Variety can absorb Variety" in a box headed "Definition", under the title "Ashby's Law of requisite Variety", with a footnote to the 1956 book.[4] We have cited it the same way.

What the law actually says

In logarithmic form the law says that the variety of outcomes cannot be pushed below the variety of disturbances less the variety of the regulator. Three details of the derivation matter for the rest of this chapter.

The first is a stipulation about the outcome table. The proof rests on it, and without it the bound weakens. Ashby's own generalization patches the gap with a constant, and in practice the constant absorbs the entire empirical content.[5]

---

[3] Beer, *Brain of the Firm*, 308, in the sentence quoted in the note on Beer's compression above: the law "can be met either by expanding regulatory variety to absorb evolutionary variety, or by curbing evolutionary learning until variety in the economy matches the regulatory variety." In *The Heart of Enterprise* the design task is put the same way: "it is management's job to DESIGN the necessary amplifiers and attenuators" (97), followed there by the First Principle of Organization. Pages read 29 September 2026 in the PDFs.

[4] Martin Pfiffner, *The Neurology of Business: Implementing the Viable System Model*, trans. Mark Kyburz (Springer, 2022), 143, https://doi.org/10.1007/978-3-031-14260-4. The box is in chapter 9, "Mastering Complexity (Excursus)," section 9.2, "Coping with Complexity"; its footnote 6 reads "W. Ross Ashby, An Introduction to Cybernetics (London: Chapman & Hall, 1956)." Page checked 29 September 2026 in the Springer PDF.

[5] Ashby, *An Introduction to Cybernetics*, 204–8. The stipulation is at section 11/5: "let us consider, then, only those tables in which no column contains a repeated outcome." Section 11/8 shows what survives when the tabular form is abandoned; section 11/9 replaces the stipulation with the assumption that each element is repeated k times in a column, and the

The second is the notation. Ashby derives the whole result in Shannon's terms and states the homology outright: a regulator's capacity as a regulator cannot exceed its capacity as a channel of communication, and "his 'noise' corresponds to our 'disturbance', his 'correction channel' to our 'regulator R'."[6] Requisite variety is channel capacity, applied to management. That gives it a lineage, and the lineage is why this law is secure where the VSM's looser mathematical claims are not. A result that reduces to information theory inherits its clarity about what is assumed and what is delivered. We say this too rarely. It is the strongest thing we can say about any formal claim we own, and it shows by contrast how little of the rest of the model has anything comparable behind it.

The third is the status of the result. Ashby is explicit: "The law states that certain events are impossible... It has nothing to do with the properties of matter."[7] A designer can use that to rule out an arrangement before building it. Anyone who wants the law to say something about a particular firm on a particular Tuesday gets nothing from it, which explains why almost nobody has tried to measure variety in a real organization. The few attempts that exist have not agreed on what should be measured. Beer saw the limit himself. Could anyone "USE this variety measure in a real-life situation," he asked, and "say: this is the variety of the chemical industry, say? Of course we could not." The use he proposed instead "is not the naive use of counting states, but of matching state generators."[8]

That last observation is about our literature, and a single counterexample would overturn it: a documented case in which two teams applied a stated counting procedure to the same organization and arrived at comparable numbers. We are not aware of one. Until one appears, the practical consequence stands. The law licenses an argument about whether a proposed arrangement could work. It does not license a number, and diagnostics that report variety figures are reporting the output of a counting convention. Chapter 12 takes up the one lineage in organization theory that did make a comparable argument operational.

---

bound becomes $V(O) \geq V(D) - \log k - V(R)$, or in entropy terms $H(D) - K - H(R)$. K is Ashby's constant, and it is unbounded by anything in the theory.

[6] Ashby, *An Introduction to Cybernetics*, 209–11. The capacity sentence ("R's capacity as a regulator cannot exceed R's capacity as a channel of communication") is at 209, section 11/11; the Shannon homology at 211. Ashby's reference is to Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* 27, no. 3 (1948): 379–423, and 27, no. 4 (1948): 623–56, https://doi.org/10.1002/j.1538-7305.1948.tb01338.x; the correction-channel result Ashby calls Theorem 10 is in Part II, section 12, at 409.

[7] Ashby, *An Introduction to Cybernetics*, 208 (section 11/10). Unelided: "The law states that certain events are impossible. It is important that we should be clear as to the origin of the impossibility. Thus, what has the statement to fear from experiment? It has nothing to do with the properties of matter."

[8] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 86–87. The question and the answer are on page 86, the proposal on page 87.

Who counts the states

Ashby introduces variety with a semaphore. It has two arms and 64 positions, but an observer standing far enough away can tell only 36 of them apart, and for that observer the variety is 36. The lesson he draws is general: "a set's variety is not an intrinsic property of the set: the observer and his powers of discrimination may have to be specified if the variety is to be well defined."[9] Two years later he put it more plainly: "If two observers differ in the distinctions they can make, then they will differ in their estimates of the variety."[10]

Variety is a count of distinguishable states, and nothing in the law fixes who does the distinguishing. Change the partition and the count changes; change the count and the law's verdict changes with it. The observer enters before the arithmetic begins. Ashby saw this and said so, but he did not close the gap. Beer kept the observer. To state a system's variety, he wrote in 1966, everything depends on "how the system is defined," and so "on who defines it"; and the system is itself a choice: "A system is not something given in nature, but something defined by intelligence," and two people "will not necessarily agree on the existence, or nature, or boundaries of any system so detected."[11] What went missing when the law was passed on is the consequence, in the quotations we use and in our practice, along with the stipulation and the constant.

Care will not repair this, and it is not a weakness peculiar to Ashby. Any formalism that starts by counting states has the same opening. What matters for us is that the same choice turns up twice more in the model, and there too we have treated it as a technical detail. Someone draws a system boundary and decides what is inside and what is environment; Chapter 5 shows how much of the recursion argument hangs on that decision. Someone decides what a transducer will preserve as a message crosses a boundary, which is the subject of Chapter 25. Counting states, drawing boundaries, specifying transducers: one political choice, made in three places. Beer connected the first two and grounded both in purpose as someone perceives it. What separates one state of a system from

[9] Ashby, *An Introduction to Cybernetics*, 125 (section 7/5). The example that precedes it is the two-armed semaphore: 64 arm positions, of which a distant observer can distinguish only 36, "so to the distant observer … the variety is 36, not 64."

[10] W. Ross Ashby, "Requisite Variety and Its Implications for the Control of Complex Systems," *Cybernetica* 1, no. 2 (1958): 83–99, in the opening section "Variety." The sentence is parenthetical in the original, and Ashby adds that he will "not, however, have to treat this complication." We read the retyped Principia Cybernetica version, which does not carry the journal pagination; page to be taken from the journal.

[11] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966; Stafford Beer Classic Library reprint, 1994), 242–43 and 247–52, pages as in the 1994 reprint. The boundary sentences are at 242 and 243, where Beer adds that "it is possible to say only that the treatment of a certain collection of things as a system is helpful," to be tested "by checking on the predictive capability of the insight." The variety sentence is at 247, where Beer warns that "the lessons of the last section are already being forgotten"; pages 247–52 then show seven things whose variety is 7, 2, 21, 42 or 2^42, depending on what they are taken to be elements of. Pages read 29 September 2026 in the PDF.

another, he wrote, is "whether the resulting change of state serves, or has no bearing on, the purposes of the system," and "systemic purpose (as we saw) is a subjective phenomenon."[12] He did not connect the transducer, and neither he nor we supplied a procedure for settling a disagreement between observers.

Naming the choice does not close it. It does stop us from presenting a variety calculation, a boundary or a transducer specification as a finding when it is a decision. And it tells us where to look when two competent analysts, using the same method on the same organization, produce diagnoses that cannot both be right.

The theorem that is not a theorem

Take a regulator that holds its system at one outcome, always the same one, and let that outcome be a catastrophe. There is no variation and no surprise. By the criterion of the paper that System Four, and System Five's models of Systems Three and Four, stand on, this regulator is optimal.

The paper is Roger Conant and Ross Ashby's, 1970, and its title states the claim: every good regulator of a system must be a model of that system.[13] We invoke it constantly and read it rarely. Beer, on the same page of *Brain of the Firm* as the *absorb* sentence, recalls it as "the cybernetic theorem that declares a regulator to be effective only insofar as its model of what is regulated is adequate."[14] The weight in Beer's sentence rests on *adequate*. The theorem's measure of success is the entropy of the outcome set, to be minimized. The paper defines a set of good outcomes at the start and then

[12]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 98; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 158. The sentence at 158 continues: "rather than a property of the system independent of its instigators, participants, and observers." It is a step in the argument whose conclusion, on the same page, is that "Freedom is in principle a computable function of systemic purpose as perceived." Beer also links his account of systems to his account of models: it "is highly reminiscent of the account given earlier of models" (*Decision and Control*, 243). Pages read 29 September 2026 in the PDFs.

[13]Roger C. Conant and W. Ross Ashby, "Every Good Regulator of a System Must Be a Model of That System," *International Journal of Systems Science* 1, no. 2 (1970): 89–97, https://doi.org/10.1080/00207727008920220. The optimality criterion is declared at 91: "In this paper we shall use the last measure, H(Z), and we define 'successful regulation' as equivalent to 'H(Z) is minimal'." The paper does define a set of good outcomes, "(2) The set G, a sub-set of Z, consisting of the ' good ' events, those ensured by effective regulation" (90), and names membership of a goal-set G as one criterion of success, beside a small root-mean-square and a small entropy (91). It then chooses H(Z), and G plays no part in the theorem. Nothing in the theorem distinguishes a low-entropy outcome set that is good from a low-entropy outcome set that is fatal. Baez noticed the same gap: "Where did the subset … of 'good' outcomes go? Shouldn't that play a role?" (see the note on Baez below).

[14]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 308. The sentence opens: "Remembering the cybernetic theorem that declares a regulator to be effective only insofar as its model of what is regulated is adequate, we see in monetarism a diminution in variety of the real economic world entailed by a regulatory model that cannot encompass more." The next sentence is the one quoted in the note on Beer's compression above. Beer does not name Conant and Ashby here. Page checked 29 September 2026 in the scan of the Wiley second edition.

does not use it again. On that measure the catastrophic regulator scores zero, which is perfect. A good regulator in Conant and Ashby's sense is a *reliable* one; whether it is a desirable one lies outside the mathematics. For an organization this is the sharpest point in the paper. The theorem can support System Four's need for a model only after someone has shown, on other grounds, that the set-point being defended is worth defending.

Beer leaned on the paper well beyond that sentence. In *The Heart of Enterprise* the principle "Every regulator must contain a model of that which is regulated" is "precisely an explanation of the necessity for System Four," and "the Conant-Ashby Theorem" returns when System Five has to hold models of Systems Three and Four, in the argument that leads to his Law of Cohesion. In a later lecture he wrote that "the formal proof" of the claim that a regulator is effective only if it "contains a model that exactly reflects the complexity of the system intended" is "available in a theorem devised by Conant and Ashby," and that "because of the Conant/Ashby theorem, any other approach is doomed."[15] Each time the theorem arrives as a statement about variety: the model "must have requisite variety," and "if this model has a given level of variety, it will handle only that equivalent variety in the outside world." Whether the 1970 proof delivers that reading is the question for the rest of this section.

The authors' own concessions are in the text, in the comments that follow the proof. On the universal claim of the title: "Although not all optimal regulators are models of their regulands, the ones which are not are all unnecessarily complex." On their own account, then, some good regulators are not models. On the proof: it, "by avoiding all mention of the inputs to the regulator R and its opponent S, leaves open the question of how R, S, and Z are interrelated." There is no sensing in it, no channel, no noise and no dynamics. On the environment: the assumption that the disturbance distribution is constant can be relaxed only to systems where it changes "slowly with time". That excludes the non-stationary environments that are System Four's whole reason for existing.[16]

[15] Beer, *The Heart of Enterprise*, 234–35 and 352. At 234 the principle is set as a displayed statement, followed by "This is precisely an explanation of the necessity for System Four"; the variety sentence is at 235. At 352 Beer names "the Conant-Ashby Theorem" and continues: "Every regulator contains a model of whatever is regulated. This model must have requisite variety"; "System Five, then, must necessarily contain models of System Three and Four." The Law of Cohesion is stated at 355. Stafford Beer, *Think Before You Think: Social Complexity and Knowledge of Knowing*, ed. David Whittaker (Wavestone Press, 2009), 110–11, in "The Will of the People": the "formal proof" sentence is at 110, "any other approach is doomed" at 111. Pages read 29 September 2026 in the PDFs.

[16] Conant and Ashby, "Every Good Regulator," 96. All three passages are in the four comments that follow the proof ("First… Second… Third… Last"): "although not all optimal regulators are models of their regulands, the ones which are not are all unnecessarily complex"; "the proof of the theorem, by avoiding all mention of the inputs to the regulator R and its opponent S, leaves open the question of how R, S, and Z are interrelated"; and, on relaxing the constancy of p(S), "if the statistics of S change slowly with time, the theorem holds over any period throughout which p(S) is essentially constant."

The second concession should trouble a community whose model is a diagram of channels. A proof that mentions no inputs, no channel and no noise cannot warrant a claim about how an organization senses its environment. Chapter 4 shows that the model's own channels fare only a little better: Beer stated their principles, and neither he nor the practice that followed specified who owns them or how they are measured.

Others have read the paper closely. In 2016 John Baez observed that the lemma carrying the proof is close to trivial, and that the step from *the regulator's state is a function of the system's state* to *the regulator contains a model* is never argued.[17] In 2021 John Wentworth wrote that it "may be the most misleading title and summary I have ever seen on a math paper", and showed that on the definition given, the identity function qualifies as a model. His repair adds expected-utility optimality, information arriving over time, and a cost of storage, and it recovers a real result: the internal state of the minimal optimal regulator must be isomorphic to the Bayesian posterior.[18] Beer had met the identity function long before, from the other side, and welcomed it: since every regulator must be an effective model, it "ought to consist of an isomorphic mapping of the system," and "the guaranteed isomorphic mapping is the identity mapping."[19] In 2025 Virgo, Biehl, Baltieri and Capucci gave the scholarly version, in a preprint on good regulators for embodied agents. The 1970 paper, they write, "doesn't strictly succeed, even in its own terms, in showing what its title claims."[20]

[17] John Baez, "The Internal Model Principle," *Azimuth* (blog), January 27, 2016, https://johncarlosbaez.wordpress.com/2016/01/27/the-good-regulator-theorem/. The post's title and its URL slug disagree; the title above is the one the post carries. Having stated the lemma, Baez asks: "What in the world does this have to do with a good regulator containing a model of the system it's regulating? Well, I can't explain that as well as I'd like—sorry." Further down he restates Conant and Ashby's conclusion as "the state of the regulator should be completely determined by the the [sic] state of the system" and adds: "this, I believe, is what they mean by" the paper's title. On the lemma, in his own comment on the post of January 28, 2016: "the proof of this lemma is so easy that I can't imagine the theorem, whatever it is, is very helpful for actually doing practical things." Baez's wording is more tentative than our summary in the text. He does not say the step is never argued; he says he cannot explain the connection well, offers a rough version of it himself, and takes the function reading to be what the title means.

[18] John Wentworth (posting as johnswentworth), "Fixing the Good Regulator Theorem," AI Alignment Forum, February 9, 2021, https://www.alignmentforum.org/posts/Dx9LoqsEh3gHNJMDk/fixing-the-good-regulator-theorem. The judgement quoted in the text is his opening: "This may be the most misleading title and summary I have ever seen on a math paper." His repair adds expected-utility optimality, information arriving in more than one chunk and having to be stored until decision time, and a cost on what is stored; the recovered result is that the model "is isomorphic to the Bayesian posterior."

[19] Beer, *Think Before You Think*, 52, in "An Open Letter to Heinz von Foerster." The passage opens: "Since every regulator must be an effective model of the system regulated, it ought to consist of an isomorphic mapping of the system," and continues, "it is evident but not trivial that the guaranteed isomorphic mapping is the identity mapping." Beer draws the opposite moral from Wentworth: he proposes to approach an identity mapping of a large system with a network of microprocessors. Page read 29 September 2026 in the PDF.

[20] Nathaniel Virgo, Martin Biehl, Manuel Baltieri, and Matteo Capucci, "A 'Good Regulator Theorem' for Embodied Agents," arXiv preprint arXiv:2508.06326 (August 4, 2025), https://doi.org/10.48550/arXiv.2508.06326. Published as "A 'Good Regulator Theorem' for Embodied Agents," in *ALIFE 2025: Ciphers of Life: Proceedings of the Artificial Life Conference 2025*

Their reconstruction replaces *being* a model with *having* one. Through an interpretation map, an observer may attribute belief states to a regulator, provided they update consistently with the system's dynamics. The price is an admission that the attributed model may be trivial.

From a complaint into a claim

That price is the useful part.

Chapter 2 follows how our teaching reduced System Four to environmental scanning. A practical failure often goes with that picture: a System Four whose scanning nobody's decisions depend on. The function exists, it produces material, and nothing downstream moves. We have always described this as a complaint, and complaints of this kind are hard to settle. The people doing the scanning point to their output. The people ignoring it point to their workload. Neither side can be shown to be wrong, because there is no shared standard to measure them against.

The reconstruction supplies one. The complaint becomes a claim: the model attributable to this regulator is trivial. Someone can dispute that with evidence, because it says something specific. No assignment of belief states to this regulator does any work in predicting what it does. To our knowledge nobody has yet applied the reconstruction to an organization, and we do not claim otherwise. What it offers is a familiar grievance turned into a proposition people can argue about, which is more than the 1970 theorem ever gave us.

The result Conant and Ashby are believed to have proved does exist, and it belongs to someone else. Francis and Wonham's internal model principle, 1976, states that a controller achieving asymptotic regulation robust to parameter perturbation must contain a copy of the dynamics generating the exogenous signals. It has a clean scope, it is taught in graduate control, and half a century of engineering stands behind it. The load under System Four is real, and Francis and Wonham carry it.[21,22]

(MIT Press, 2025), paper 46, https://doi.org/10.1162/isal.a.874; we cite the preprint because that is the version we read. Version 2 (21 August 2025) keeps the quoted sentence unchanged (checked 26 September 2026). In the original the quoted sentence reads "(C&A) doesn't strictly succeed, even in its own terms, in showing what its title claims," where (C&A) is the paper's own abbreviation for Conant and Ashby. The authors credit Baez and Wentworth with having made the point first.

[21] Bruce A. Francis and W. Murray Wonham, "The Internal Model Principle of Control Theory," *Automatica* 12, no. 5 (1976): 457–65, https://doi.org/10.1016/0005-1098(76)90006-6.

[22] A second route from the law to the need for a model, which does not pass through Conant and Ashby, is the law of requisite knowledge: having enough actions is not sufficient, since the regulator must also be able to select the one that fits the disturbance. In entropy form, $H(E) \geq H(D) + H(R|D) - H(R) - K$. E stands for the essential variables, the ones the regulator has to keep within bounds; K, as in Ashby, is the variety removed by passive buffering; and $H(R|D)$ measures the regulator's uncertainty about which action to take given the disturbance. Heylighen and Joslyn credit this extension to Aulin: "Aulin has shown that the law of requisite variety (4) can be extended to include knowledge or ignorance by simply adding this conditional uncertainty term." They give no reference for it. The same term appears in Aulin-Ahmavaara's 1979 paper, where H_D(R) "represents the ignorance of the regulator about

Their scope condition should stay in view, since dropped scope conditions are what this chapter is about. Francis and Wonham prove their result for a class of control problems with stated dynamics, and an organization is not obviously such a system. What their theorem gives us is a rigorous statement of the form we thought we already had, and a standard for judging whether an organizational reading of it is being argued for or merely assumed.

A derivation that went almost uncited

On 29 September 2026 OpenAlex recorded 1,065 citations for Conant and Ashby's paper and 2,879 for Francis and Wonham's. For a third paper, published in a cybernetics journal, it recorded 44.[23]

The third paper is A. Y. Aulin-Ahmavaara's law of requisite hierarchy, published in *Kybernetes* in 1979, the year *The Heart of Enterprise* appeared.[24,25] Aulin argued formally that hierarchy is what compensates when regulatory ability is limited. That is close to Beer's recursion, though not the same object, and it comes with the kind of formal derivation Beer promised and, as far as we have found, never gave.[26]

Counts from one database are a crude instrument. Coverage differs by field and by decade, citations in books are badly captured, and no count

---

how to react correctly to each appearance of a disturbance D" (260). Francis Heylighen and Cliff Joslyn, "Cybernetics and Second-Order Cybernetics," in *Encyclopedia of Physical Science & Technology*, 3rd ed., ed. R. A. Meyers, vol. 4 (Academic Press, 2001), 155–70; we read the preprint; the page range matches the Vrije Universiteit Brussel record of the chapter, and we have not seen the printed volume. We cite it as a formulation that makes the knowledge requirement explicit, not as a proof.

[23]On 29 September 2026, OpenAlex recorded 44 citations of Aulin-Ahmavaara's paper, against 1,065 for Conant and Ashby (1970) and 2,879 for Francis and Wonham (1976); on 7 August 2026 the counts were 43, 1,040 and 2,838. Citation counts from one database are a crude instrument and we use them only for the order of magnitude.

[24]A. Y. Aulin-Ahmavaara, "The Law of Requisite Hierarchy," *Kybernetes* 8, no. 4 (1979): 259–66, https://doi.org/10.1108/eb005528; the law is stated at 262. Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979).

[25]The 1979 paper is not the first statement of the law. Aulin-Ahmavaara writes that its content "seems to have been published in English for the first time in an article of the present writer published in 1975 in Cybernetica," and his reference list gives A. Y. Ahmavaara, "Cybernetics as the foundational science of action," *Cybernetica* 3 (1975): 171–200 ("The Law of Requisite Hierarchy," 262; reference 1 at 266). We have not read the 1975 paper. The coincidence of dates with *The Heart of Enterprise* holds for the *Kybernetes* paper, not for the idea, which was in print four years earlier.

[26]This identification may be overstated. Aulin-Ahmavaara's hierarchy is a hierarchy of control: governors placed above the regulators compensate for the regulators' ignorance of which action fits which disturbance. His law says that "the lack of regulatory ability can be compensated for, up to a certain amount, by a greater hierarchy in organization" (262), and he applies it to social classes and the organization of society as a whole. The paper does not mention Beer, recursion or viability. The law maps at least as naturally onto the metasystem above System One as onto recursion, the nesting of viable systems inside viable systems. Beer stated his Recursive System Theorem in 1979 and wrote that "the proof offered later will be a topological proof" (Stafford Beer, *The Heart of Enterprise*, John Wiley & Sons, 1979, 118); a full-text search of the book found no passage presented as that proof.

separates a paper that was read from a paper that was merely listed. The ratio survives all of that. A formal treatment of the structural claim at the center of the VSM, published in the same year as the book that argues for that structure, has been cited tens of times. The two control-theoretic results discussed above have been cited thousands of times. No better instrument is needed to see that we left an available argument lying where it was. Whether Aulin's derivation actually establishes recursion in the form the VSM needs is a sharper question. It would take a close reading of the paper against Beer's claim. This book does not do that reading, and we would like to see it done.

Aulin is not an isolated case. Work that would have given our claims formal support has repeatedly been done outside the tradition and gone unread inside it. Chapter 12 follows another instance through Galbraith, who arrived at a variety argument for organization design without citing Ashby at all.

What the law becomes when someone has to measure it

We have said that this tradition quotes a law and uses a vocabulary. There is one place to check that, because one researcher had to turn requisite variety into a question an informed outsider could answer.

Michael D. Pfiffner's 2017 dissertation states the law at the head of the section that introduces the relevant items. It keeps Ashby's verb and glosses his noun: "Ashby's well-known law of requisite variety states that only variety (i.e. complexity) destroys variety." On the same page, the requisite-variety items of his instrument ask whether the organization's "ability to react (i.e. the capability to react in a timely, appropriate and formally correct manner), scope of action (the set of suitable behaviours) and the range of offerings (available solutions to a problem)" are "always sufficient to master expected" internal or external challenges. A second item asks the same about unexpected ones.[27]

Nothing in these items is counted. They enumerate no states of the regulator and none of the regulated system, and they set up no inequality between the two. They ask whether the people concerned can cope, and the answer is a judgment of adequacy on a scale. Ashby's law has become a question about whether an organization seems equal to what it faces. The step can be followed on one page: variety becomes complexity in the sentence that states the law, and complexity becomes adequacy in the items below it.

[27] Michael Dominik Pfiffner, *System viability of organizations and the aetiology of organizational crisis: A Quantitative Assessment of Stafford Beer's Viable System Model* (PhD diss., Universiteit Utrecht, 2017), section 4.6.3.10, "Principle of Requisite Variety," Table 18 (Questionnaire Principle Requisite Variety), items P.3 and P.4 (Principles Repertoire), 165. The items are also reproduced in the dissertation's appendix. The statement of the law quoted in the text is the first sentence of section 4.6.3.10, on the same page as Table 18. Read in full text 14 September 2026; page checked 22 September 2026 against the Utrecht PDF, and the statement of the law on 29 September 2026 in the Utrecht repository (publisher) version.

The instrument could not have done otherwise. Counting states is not available to a survey of bankers and consultants, and it is barely available to anyone, for the reasons set out under who counts the states. But that is the finding. When a careful researcher had to make the law operational, what came out was the vocabulary, and the vocabulary is what we have been quoting all along. Chapter 18 examines the study's further use of that item, which is less benign: its near-negation appears on the other side of the same regression.

What this chapter hands forward

Stop citing the good regulator theorem as though it were proved. Cite Francis and Wonham for the rigorous claim and Virgo and colleagues for the modern reconstruction. The authors of the 1970 paper qualified its title inside the paper itself, and we went on invoking it. We take that as evidence that we have not read our own foundations closely enough. It is also the cheapest correction open to us.

Own the partition problem instead of inheriting it three times. Who counts the states, who draws the boundary, and who defines the transducer are one question wearing three coats. Better technical work will not close it. It marks the point at which an observer, and with the observer a politics, enters the model.

Receipt — one corrected citation practice; cost: free. Stop presenting Conant and Ashby's title as a proved universal. Cite Francis and Wonham for robust regulation through an internal model, and Virgo and colleagues for the modern reconstruction. This is a practice change, not an experiment: it is retired only by a proof that restores the universal claim under conditions relevant to organizational regulation.

This is the smallest receipt in the book, and deliberately so. Chapter 9 asks why so little in this tradition has ever cost us anything, and part of the answer is that we have rarely been anywhere that being wrong was expensive. Changing a citation practice costs no one anything. If we cannot make this correction, the obstacle is not resources.

Chapter 4

Channels Without Owners

In seminars and in print, we describe the algedonic↗ channel as the one that escapes ordinary hierarchy. Beer says the same: algedonic signals are "organized to *break through* recursivity."[1] He also says how a signal qualifies. A statistical filter watches chosen points of stability, splits off the readings that show incipient instability, and decides whether to alert System Five. His phrase for it is "a device, operating in real-time, that monitors chosen points of stability by statistical methods," and the metasystem sets its criteria.[2] What Beer does not say is who chooses the points of stability, who may contest the filter, and who audits it, on whose behalf. Those are the questions an institution has to answer, and neither Beer nor the practice that followed him answered them. We draw this channel as a mechanism and have never specified it as an institution, and a channel without those answers has no owner anyone could name.

Beer counts six vertical channels between the operating units and their metasystem, only one of which, he complains, is usually acknowledged,[3] and he sets the algedonic channel beside them, "orthogonal to the calming system."[4] Three of these carry this chapter: the resource bargain, audit, and the algedonic channel.

Four neighboring literatures approach the gap from outside the model. None of them was written against the VSM, and only one was written after it. Each says something specific about the relationship between a level and the level containing it: what can cross, who is entitled to send it, whose interests the crossing serves, and how large the asymmetry between the two levels may become. Taken in order they make one argument, and the last of them, about variety, is the reason the first three are more than stylistic complaints.

[1] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 407, where algedonic signals, "alone of all the signals in the enterprise," break through the recursions.

[2] Beer, *Heart of Enterprise*, 386 for the quotation, and 275, 381–82; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 178–79: "Statistical criteria must be established to ensure that ascending algedonic information … is not simply absorbed"; Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 133, where the signal passes through "its own algedonic filter" on the way to System Five.

[3] Beer, *Heart of Enterprise*, 142–43: "These six vertical loops (of which only one, the command loop, is customarily acknowledged …)"; Beer, *Diagnosing the System*, 81–82, for the six: corporate intervention, the resource bargain, operational linkages, environmental intersects, System Two, and audit; and 68 on System Two: "Its function is not to command, but to damp oscillations."

[4] Beer, *Heart of Enterprise*, 406.

Revolt, remember, and the coupling we never drew

Holling's adaptive cycle is the one member of this group that comes after Beer rather than before him: resilience in 1973, the cycle itself in 1986, panarchy proper at the turn of the century. Systems traverse exploitation, conservation, release and reorganization. The cycles nest across scales that differ in speed and extent, and the levels are coupled in both directions. In *revolt*, a fast small collapse cascades upward into a slower larger level. In *remember*, the slow large level supplies accumulated structure downward to reconstitute a small one after its collapse.[5]

Beer's channels have no counterpart to either. His vertical channels command, bargain, coordinate, inspect and alarm, and revolt and remember are contagion and inheritance. The algedonic signal is arguably a degenerate revolt, a pain report that bypasses levels. Remember has no counterpart at all, so the model has no account of how a destroyed unit is reconstituted out of the memory of its container.

The word *degenerate* is doing work there, and it needs care. A revolt in Holling's sense is addressed to no one. It propagates because the levels are materially coupled, and the larger level receives it whether or not it has a receiver. Our algedonic signal is addressed. In Beer's primary design nobody writes it, since a filter generates it, but it has a recipient and an obligation to respond: if the top did not respond, he says, the system would "*lose its cohesiveness*."[6] Beer also has authored forms, the informal telephone call to the boss and the Chilean algedonic meter, on which a person "does not have to explain anything."[7] The addressing makes the channel governable. It also makes the specification questions unavoidable, because an authored signal can be withheld or forged, and a filter's criterion can be set by whoever owns it. Remember is harder still. The containing system would have to hold structure on behalf of a unit that no longer exists, and be able to reissue it into a rebuilt unit with no continuity of personnel with the old one: the operating model, the local knowledge, the relationships, the standing to act.

We do not claim the model forbids this. System Four's model of the environment and System Five's account of identity might in principle be where such structure is held, and practitioners plainly do reconstitute failed units in the field. The claim is narrower and, we think, harder to dismiss: nothing in the canonical description tells us where reconstitution capacity

[5] C. S. Holling, "Resilience and Stability of Ecological Systems," *Annual Review of Ecology and Systematics* 4 (1973): 1–23; C. S. Holling, "The Resilience of Terrestrial Ecosystems: Local Surprise and Global Change," in *Sustainable Development of the Biosphere*, ed. William C. Clark and R. E. Munn (Cambridge University Press, 1986), 292–317, where the adaptive cycle first appears; Lance H. Gunderson and C. S. Holling, eds., *Panarchy: Understanding Transformations in Human and Natural Systems* (Island Press, 2002), and C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405, for revolt and remember.

[6] Beer, *Heart of Enterprise*, 407.

[7] Beer, *Diagnosing the System*, 133, for the informal route; Beer, *Brain of the Firm*, 283–84, for the meter.

lives, who is responsible for maintaining it, or how one would know it had decayed. Cases would settle it, where argument cannot. Take a set of documented instances in which an operational unit was destroyed and later rebuilt inside the same containing system. Ask what supplied the structure, whether the container had retained it deliberately, and whether anyone had ever been accountable for retaining it. Chapter 5 takes up the related question of what fixes the levels in the first place.

Whose audit is it

Ostrom arrives at the same boundary from the other side, and with evidence.

Read her eight design principles for long-enduring common-pool resource institutions with the VSM in hand. Clearly defined boundaries. Congruence between the appropriation and provision rules and local conditions. Collective-choice arrangements, in which most of those affected by the operational rules can participate in changing them. Monitoring by monitors who "are accountable to the appropriators or are the appropriators". Graduated sanctions. Rapid, low-cost, local conflict resolution. Minimal recognition by external authorities of the right to organize. And, for resources embedded in larger systems, nested enterprises, with appropriation, provision, monitoring, enforcement, conflict resolution and governance "organized in multiple layers".[8]

The fourth principle is the decisive one here. It resembles the audit channel (the System Three-star of *Diagnosing the System*) and goes further than Beer did. Beer's audit is sporadic, and it reaches the operations "by agreement with System One management":[9] a consent condition, held by the unit's managers. Ostrom's condition is held by the people whose conduct and resources are being inspected. Its monitors must be accountable to them, or be drawn from them. A channel's independence is a property of who owns it, not of the box it bypasses. The seventh principle, and Chile in September 1973, belong to the Cybersyn argument in Chapter 8. The eighth, nested enterprises, is the ladder question and belongs to Chapter 5. We note both handoffs and reopen neither.

The fourth principle transfers directly to the algedonic channel, which is why it appears here and not only in Chapter 14. If a Three-star audit is not independent simply because it routes around the line, then an algedonic signal is not a report from the affected simply because it routes around the line either. In the implementations we have seen described (we have

[8] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), 90 (table 3.1). Principle 4 reads in full: "Monitors, who actively audit CPR conditions and appropriator behavior, are accountable to the appropriators or are the appropriators." Principle 8 reads: "Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organized in multiple layers of nested enterprises," and Ostrom states it only "for CPRs that are parts of larger systems."

[9] Beer, *Diagnosing the System*, 82. The label "3*" for the audit channel is from this book; in* The Heart of Enterprise* the same label names a model inside System Five (352).

not surveyed them systematically), the entitled sender has usually been a manager, reporting on the condition of the unit he manages. Beer's own design points the same way. The filter that decides whether pain reaches the top is designed by the metasystem,[10] the party the signal is meant to wake. Ostrom's fourth principle exists to prevent exactly that structure. We adopted the channel with the ownership Beer gave it, and did not ask whether it was the right owner.

This evidence has a limit. Ostrom's principles were induced from common-pool resource institutions, which have appropriators with standing claims and no command channel. Firms and public agencies have both a command channel and employees who can be dismissed. Whether the ownership requirement survives that difference, or has to be rebuilt in some other form, is an open question. Chapter 14 sets out what a comparative study would have to look like.

Who is at the boundary

Critical systems heuristics asks twelve boundary questions, each in an *is* and an *ought* mode, about who is served, who decides, what counts as knowledge and what legitimates a design; Chapter 16 sets them out in full.[11] Three bear on this chapter: the measure of improvement, the guarantor of success, and the tenth, which asks who is witness to the interests of those affected but not involved.

Applied to us, the questions are uncomfortable in the right way. Who names the identity in the step Beer left to the eye of the beholder?[12] Who fixes the recursion levels? Whose measure of improvement do the System One indices encode? And what is the assumed guarantor of success: viability itself, or the model's own theorems? Both are internal to the designer's frame, and so they can justify nothing to anyone excluded from it.

The tenth question is the one that bears on the algedonic channel. The channel has no seat for the person who is affected and was never a member.

[10] Beer, *Heart of Enterprise*, 386, where the limits of each point of stability are "determined by the criteria laid down by the metasystem"; and 382, where the resolution filter on a manager's input "has been designed by the metasystem of which our manager is a member." The resolution filter "does not know whether John or Mary is operating that machine today"; the manager's perceptual filter, by contrast, is designed by the manager himself.

[11] Werner Ulrich, *Critical Heuristics of Social Planning: A New Approach to Practical Philosophy* (Paul Haupt, 1983); Werner Ulrich and Martin Reynolds, "Critical Systems Heuristics," in *Systems Approaches to Managing Change: A Practical Guide*, ed. Martin Reynolds and Sue Holwell (Springer, 2010), 243–92. The ordering used here is Ulrich's own: the twelve questions fall in four groups of three — sources of motivation, of power, of knowledge, of legitimation — and each is asked in an *is* and an *ought* mode. The measure of improvement is question 3, the guarantor question 9, and the witness to those affected but not involved question 10, in the numbering of Ulrich's checklist of boundary questions as reproduced in Werner Ulrich, "A Brief Introduction to Critical Systems Heuristics (CSH)" (ECOSENSUS project, Open University, 2005), 11 (fig. 4), which he dates to 1987.

[12] Beer, *Heart of Enterprise*, 9: "we simply cannot attribute purposes, or even boundaries, to systems as if these were objective facts of nature. The facts about the system are in the eye of the beholder."

In Cybersyn, as Chapter 8 shows, an algedonic signal was more than a cry of pain; it was a transfer of control upward. The architectural consequence is that the person who bears the pain may neither author the signal nor control what the escalation does. Beer's algedonic meter is the one design in which that person authors it, and it stayed outside the architecture as we teach it.

This is Ostrom's fourth principle again, in a wider frame. Ostrom asks whether the monitor is drawn from those monitored; Ulrich asks whether anyone at all speaks for those outside the membership boundary. A specification that answered the first without the second would still leave the channel closed to the neighbor, the patient, the downstream community.

Nothing in critical systems heuristics is incompatible with the Viable System Model. The two operate at different logical levels: Ulrich interrogates the boundary, Beer designs inside it. In intent, System Five and the algedonic channel are a witness function. The model has the slot, and Beer says who should fill it: System Five "consists of many people," possibly "simply workers," in Chile *el pueblo*, and later all stakeholders.[13] He gives no procedure for seating them, and nothing for those who are affected but outside the boundary. That defect can be fixed without refuting anything, and fixing it is the cheapest high-value reform available to us. Chapter 16 takes Ulrich's objections in full, including those that do not resolve so comfortably.

A regulator with no stated upper bound

The three findings so far concern who may use a channel. The fourth concerns how far apart the two ends of a channel are allowed to drift.

Ashby's law states a minimum variety for a regulator.[14] Nothing in the sources we searched for this chapter, Ashby's *Introduction to Cybernetics* and five of Beer's books, states a maximum compatible with that regulator remaining auditable by the regulated.[15]

[13]Beer, *Heart of Enterprise*, 263–64; Beer, *Brain of the Firm*, 258, and 288: "System Five within a people's government cannot be an élite ruling class"; Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 149: "All stakeholders are represented in System Five."

[14]W. Ross Ashby, *An Introduction to Cybernetics* (Chapman & Hall, 1956), 206–7 (sections 11/6–11/7): if the variety in the outcomes is to be reduced to an assigned number, "R's variety must be increased to at least the appropriate minimum." The one maximum Ashby states runs the other way. A regulator fixed in its channel capacity meets "an absolute limit to the amount of regulation (or control) that can be achieved by R" (245, section 13/2): a ceiling on what the regulator can do, not on how much variety it may hold. Chapter 3 sets out the law's conditions.

[15]The texts searched were Ashby, *Introduction to Cybernetics*, read in full for this question, and five of Beer's books, searched in full text for any stated upper limit on a regulator's variety: *Decision and Control* (John Wiley & Sons, 1966); *Platform for Change* (John Wiley & Sons, 1975); *The Heart of Enterprise*; *Brain of the Firm*, 2nd ed.; and *Diagnosing the System*. The nearest statements are limits of capacity, not of auditability. Beer writes that the maximum variety a senior management "can handle for each subsidiary is their own total capacity divided by the number of subsidiaries in System One" (*Diagnosing the System*, 51), and he reports

Beer comes closest. The metasystem "should make only that degree of intervention that is required to maintain cohesiveness," and freedom is "in principle a computable function of systemic purpose as perceived."[16] Metasystemic effort beyond what cohesion needs is, in his diagnosis, a pathology.[17] That is a ceiling on intervention, set by cohesion and purpose. It is not a ceiling on the variety the regulator may hold, and Beer gives no procedure by which the regulated could check either one.

We flag this as a scoped negative finding. It describes those texts and is not the result of a systematic literature search; a single citation we have missed would overturn it, and we would welcome one. Chapter 3 sets out what Ashby's law does and does not say. The absence matters here because our other three findings all assume that the regulated can, in principle, check the channel that speaks about them. A regulator whose variety may grow without stated bound can satisfy the law and still be unauditable in practice. The ownership requirement and the witness seat then become formalities. If there is an upper bound, we should be able to say what it is a function of.

The transducer, the mechanism that decides which differences survive passage through a channel, is the part of the variety logic its rediscoverers in Chapter 12 left alone. Part Four returns to it, in Chapter 25. A specification of who may send an algedonic signal is incomplete without an account of what the channel does to the signal in transit. Beer states the requirement: "the variety of the transducer must be at least equivalent to the variety of the channel," and channel capacity and transduction capacity are two different things.[18] He gives no way to measure either at a real boundary.

What it would cost to find out

Receipt — red-team the algedonic channel. The reform this chapter argues for begins with a document we have not written. Write a qualifying-signal specification naming entitled senders, thresholds, evidence, routing, response time, appeal and audit owner. Then give a red team the same operating data and ask it to manufacture a signal that passes while misrepresenting the condition. Price: one facilitated design session and one red-team day. If false signals pass trivially, the negative result is publishable: the channel transfers control without a defensible rule for when transfer is warranted. Chapter 9 explains why so few tests of this kind exist in our literature, and the explanation is not that they are expensive.

Bremermann's physical limit on the information-processing capacity of matter (*Brain of the Firm*, 48). The neighboring literatures of this chapter (Holling, Ostrom, Ulrich) were not searched for the question.

[16] Beer, *Heart of Enterprise*, 157–58.

[17] Beer, *Heart of Enterprise*, 411–12, on the metasystem that maintains itself as "The Establishment." The Law of Cohesion (355) is stated as an equality, not as a minimum.

[18] Beer, *Heart of Enterprise*, 101; Beer, *Diagnosing the System*, 47, where "channel capacity and transduction capacity … are not at all the same thing."

Chapter 5

The Ladder Without a Ruler

When a client asks why the branch is a recursion level and the team within it is not, we answer with an argument about identity. The argument is usually a good one. It is still an argument rather than a reading.

Beer offers a good criterion for recursion: the candidate system must be viable in its own right and one of the units that do what the containing system does; Chapter 10 states it in full, with Beer's own fences.[1] Around the criterion he supplied more than nothing and less than a ruler. There is a heuristic: the operational units are the ones that between them produce the enterprise, and if you count more than seven or eight of them, "you may be missing a whole level of recursion." There is a workbook step, "Choose and clearly define the System-in-focus," and there is, in principle, a test of predictive usefulness.[2] None of it tells two competent analysts how to settle a tie between two drawings that both pass. Beer granted as much: two observers "will not necessarily agree on ... the boundaries of any system."[3] He also chose the mathematical term *recursion*, he wrote, because it reminds us "we are not talking loosely about any kind of system contained inside another — but about an absolutely precise definition of viability," and he held that the levels answer "to mathematical criteria."[4] The precision is asserted and never supplied.

The gap is not small, and we cannot blame it on anyone else's neglect. We draw levels and call the drawing a diagnosis. Nothing in our practice tells us in advance what result a second analyst would get, and, as Chapter

[1]The formulation is ours, not a quotation. Its substance is in Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), chap. 5, 116–20: operational elements are those which "exist to do what the system does" (116), and each is a viable system that "could in principle be hived-off as a separate business" (118). Page numbers corrected 29 Sept 2026 against the calibre PDF.

[2]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 8 (the units that between them produce the company), 16 (the System-in-focus), 19 ("up to seven or eight: not many more, I hope, or else you may be missing a whole level of recursion"); Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113 and 243, on judging a model, and a recognized system, by its predictive capability.

[3]Beer, *Decision and Control*, 243: "Two people will not necessarily agree on the existence, or nature, or boundaries of any system so detected."

[4]Beer, *Diagnosing the System*, 2; Beer, *Heart of Enterprise*, 308. Two pages later he promises the transformation rules between levels "in terms of structural isomorphism" (310), and on the next he calls both the location of the origin and the extent of the recursion on either side of it "quite arbitrary": "The focus of our recursions is wherever our interest lies" (311).

9 sets out, nothing in our commercial arrangements has ever made it costly for us to find out. We have a ladder and no ruler.

Four neighboring literatures each supply something the VSM↗ lacks at this point, and each lacks something the VSM has. Set beside Beer's criterion, they show what an instrument would have to do. A fifth body of work, Ostrom's, says something less comfortable: in the one large evidence base we have habitually cited in support of recursion, recursion is the weakest thing in it.

Koestler supplied the noun

The Ghost in the Machine appeared in 1967. Koestler named the thing: a holon, the unit that is a whole to what it contains and a part to what contains it. "Parts and wholes in an absolute sense do not exist in the domains of life."[5] Beer's nouns are relative. He says "System-in-focus," or calls the contained systems "recursions of the viable system," always measured from a focus someone has chosen.[6] Koestler gave one term for the unit's double role, which in our experience is one reason VSM diagrams are harder to read than holonic ones. Naming was Koestler's contribution, and it was not a small one.

Naming has its limits, and they need stating as plainly as the achievement. A noun makes a structure discussable; it does not make it detectable. Koestler gave us a way to talk about the recursing unit without circumlocution, and Chapter 11 traces how that vocabulary reached us alongside Simon's and Jaques's. He did not give us a way to walk into an organization and find out how many holons are in it. What we can say with confidence is narrower: the relative nouns in Beer's presentation impose a cost on readers, and the holonic literature does not pay it.

PROSA turned the word into engineering

The word also acquired an engineering literature, and that literature is the most useful thing in this chapter. Holonic manufacturing systems took Koestler's word into factory automation in the nineteen-nineties. In 1998 Van Brussel and colleagues published PROSA, a reference architecture with a fixed small vocabulary of holon types: product, resource, order, and an advisory staff holon whose failure degrades the holarchy to a heterarchy rather than killing it. That last is a design idea Beer does not have; the

[5]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), 48, where the holon is coined. The sentence quoted is proposition 1.3 in both of Koestler's proposition sets, and we give it in the form of the 1969 Alpbach appendix. The 1967 book prints it in its own Appendix I, 341, with "domain" for "domains" (read in the Internet Archive scan of the Hutchinson first edition on 2 October 2026; proposition 1.3 checked against the page image).

[6]Beer, *Diagnosing the System*, 2: the embedded systems "are themselves viable systems. They are RECURSIONS OF THE VIABLE SYSTEM"; 19 for "System-in-focus." The recursive system theorem makes the viable system itself the two-sided term: "any viable system contains, and is contained in, a viable system" (Beer, *Heart of Enterprise*, 118).

VSM's metasystem↗ is not optional in the same way.[7] PROSA has been cited some fourteen hundred times.[8]

Two separate features of that architecture concern us. The first is the fixed small vocabulary. PROSA does not invite each practitioner to decide what the holon types are. The types are given, and a design either uses them or is not a PROSA design. That constraint makes independent implementations comparable, and comparability is the precondition for accumulating evidence at all. Our own practice runs the other way: the five systems are fixed, but what counts as an instance of System One↗ at a given level is settled case by case.

The second is graceful degradation. A staff holon that is by construction only an advisor can be removed, and the execution system continues in a heterarchical mode: less predictable, less well optimized, still running. We make no equivalent claim. Our metasystem is constitutive rather than advisory, so in our account a metasystem failure is a loss of viability, not a degradation. Whether this is a real difference between the two architectures or a difference in how each states its commitments is worth settling, and a comparison of documented implementations could settle it. Chapter 15 takes the thirty-year manufacturing record at length. Here it is enough that a neighboring literature made a design commitment about partial failure that we have not made.

Jaques supplied a ruler

Jaques claims that the level of work in a role is measurable by a single quantity, the time-span of discretion: the longest period a subordinate exercises discretion before the results become reviewable. He also claims that organizational strata are discrete, with boundaries at roughly three months, one year, two years, five, ten and twenty.[9] Beer's criterion for a recursion level produces a judgment about a system's identity. Jaques's

[7] Hendrik Van Brussel, Jo Wyns, Paul Valckenaers, Luc Bongaerts, and Patrick Peeters, "Reference Architecture for Holonic Manufacturing Systems: PROSA," *Computers in Industry* 37, no. 3 (1998): 255–74. The graceful-degradation property is stated most plainly not in the original paper but in Paul Valckenaers, "Perspective on Holonic Manufacturing Systems: PROSA Becomes ARTI," *Computers in Industry* 120 (2020): 103226: "To qualify as a staff holon (e.g. a scheduler), the staff holon must be no more than an advisor. Removal or disabling of a staff holon must not break the execution system composed of order holons, product holons and resource holons. Performance may degrade, performance predictability may suffer, but the heterarchical operating mode remains operational."

[8] Citation counts on 29 September 2026: 1,465 in Semantic Scholar and 1,412 in OpenAlex. Counts differ by index and drift upward; "some fourteen hundred" is the honest form of the number and should not be quoted to more precision than that.

[9] Elliott Jaques, *Measurement of Responsibility: A Study of Work, Payment, and Individual Capacity* (Tavistock Publications, 1956), for the derivation of time-span of discretion, and Jaques, *Requisite Organization: The CEO's Guide to Creative Structure and Leadership* (Cason Hall, 1989), for the stratum boundaries at three months, one year, two years, five, ten and twenty years in their settled form. The lowest five of those boundaries are the ones Rowbottom and Billis tested in 1977, a study Chapter 11 discusses; the twenty-year boundary is not independently corroborated in anything we found.

produces a number attached to a role. That is the whole of the difference, and it is enormous. Two analysts who disagree about whether a unit is viable in its own right and one of the units that do what the containing system does have no procedure for settling it, and Beer's rules of thumb do not provide one.

The limit deserves the same volume as the contrast. The validity of Jaques's ruler is contested, and we do not endorse it here. Whether time-span measures level of work or something correlated with it, whether the strata are genuinely discrete or an artifact of the measuring convention, and whether the upper boundaries have any support at all are open questions in the literature that grew around requisite organization. From Jaques we take the form of the answer, and leave the answer itself open: a single quantity, defined operationally, that two analysts can independently obtain from the same role and then compare. A contested measurement is a different kind of object from an uncontested judgment call, because a contested measurement can be tested and a judgment call cannot.

Panarchy shows what the diagram cannot hold

The contrast with Beer is the sharpest in this chapter. His recursion is structurally invariant: the same five systems at every level, always. Beer does describe modes of behavior, among them growth, retrenchment and crisis, and a cohesion that varies with the system's purpose.[10] He does not describe phases coupled across levels, and the diagram has no notation for either. Panarchy's nesting is temporal. Each level is somewhere in a cycle, different levels are in different phases at once, and the coupling between them changes with phase.[11] A System One in release does not need the System Three↗ that a System One in conservation needs. A VSM diagram cannot represent that, and every restructuring produces exactly that situation.

The limit reaches past the diagram to our ability to use the alternative. Holling's phases are considerably easier to identify after the fact than in the present tense, and coupling that varies with phase is actionable only if the phase of each level can be read while the work is going on. We are not aware of a procedure that does this reliably for organizations. The finding is therefore asymmetric. Panarchy identifies a real representational limit in our diagrams without yet supplying the instrument that would remove it. What it does establish is that the invariance we treat as a strength (five systems, every level, always) is at least in part a modeling choice. Beer

[10] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 233–35, where the algedonic controller "determines the *mode* of behaviour"; Beer, *Heart of Enterprise*, 158: "Viable systems of concentrated purpose will be closely-knit, highly cohesive. Viable systems of general purpose will be more loosely coherent."

[11] Lance H. Gunderson and C. S. Holling, eds., *Panarchy: Understanding Transformations in Human and Natural Systems* (Island Press, 2002); C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405. Chapter 4 sets out the cycle's four phases and the couplings between levels.

himself called the choice of five "somewhat arbitrary," later declared the five-fold system closed "in logic," and announced a topological proof of the invariance that we have not found stated as such.[12] Choice or discovery, the invariance costs us the ability to say that two levels of the same organization are in different regimes at the same moment. Chapter 4 makes the neighboring point about channels: what the notation cannot express tends to drop out of the diagnosis.

The evidence base we counted on says otherwise

Ostrom's evidence base yields one finding that bears directly on recursion, and it is uncomfortable. The *least* supported of the principles is the nested one, the principle closest to recursion. Nested enterprises came out at the bottom of the eleven in the 2010 review: alone at the bottom, still within the range the reviewers call moderately supportive, but the only one of the eleven whose association with success failed to reach significance at the five per cent level. Proportionality between what a member puts into a commons and what they take out came out near the top, and it is also one of the two conditions the 2016 configural analysis found holding across all three sectors. Nesting held nowhere in particular.[13]

We had counted this literature as support, and at this point it counts against us. Ostrom's multi-level organization is an independent arrival

[12] Beer, *Brain of the Firm*, 98; Beer, *Diagnosing the System*, 130: "There could not be a 'System Six' — because the five-fold system is a *closed* system, in logic"; Beer, *Heart of Enterprise*, 118, for the promised topological proof.

[13] The principle is Ostrom's eighth, nested enterprises: Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), 90 (table 3.1). Michael Cox, Gwen Arnold, and Sergio Villamayor-Tomás, "A Review of Design Principles for Community-Based Natural Resource Management," *Ecology and Society* 15, no. 4 (2010): art. 38, https://doi.org/10.5751/ES-03704-150438; and, for the 2016 configural analysis, Jacopo A. Baggio, Allain J. Barnett, Irene Perez-Ibarra, Ute Brady, Elicia Ratajczyk, Nathan Rollins, Cathy Rubiños, et al., "Explaining Success and Failure in the Commons: The Configural Nature of Ostrom's Institutional Design Principles," *International Journal of the Commons* 10, no. 2 (2016): 417–39, https://doi.org/10.18352/ijc.634. The per-principle support ratios reported by the 2010 review are quoted inconsistently in the secondary literature, so we give them here from table 3 of the paper itself, as ratios of supportive to unsupportive cases: 4B, accountability of monitors, 11.7; 2B, proportionality, 10.0; 1A, user boundaries, 5.3; 2A, congruence with local conditions, 5.0; 4A, presence of monitors, 3.8; 6, conflict resolution, 3.2; 3, collective choice, 2.8; 5, graduated sanctions, 2.3; 1B, resource boundaries, and 7, right to organize, 2.2 each; and 8, nested enterprises, 2.0. Nesting is therefore alone at the bottom and not, as we had it in an earlier draft, tied with the right to organize. It is also the only one of the eleven for which Fisher's exact test failed at the five per cent level ($P = 0.065$; the authors report it as significant at the 10 per cent level), and it carries the smallest effect size (phi = 0.262). The authors' own bands treat ratios of 2 to under 5 as moderately well supported, so nesting sits at the very edge of the weakest band they use. In the 2016 analysis nested enterprises co-occur less with the other principles in successful cases (Baggio et al., 426, 428), and the successful irrigation and forestry cases include a solution set that contains every principle except nesting (429).

at something like recursion. It is also the principle her own literature supports least: support ratio 2.0, Fisher's exact P = 0.065 and phi = 0.262.[14]

Reading the result fairly takes care in both directions. It does not refute recursion. Nesting remains inside the band the reviewers themselves call moderately supportive. The review is a synthesis of case studies, with all the selection and coding problems that implies, and a principle can be real and still fail to show a strong association with success in a body of cases where it is unevenly present and inconsistently coded. Nor does the result transfer cleanly: common-pool resource institutions are not firms, and the nesting Ostrom's coders looked for is not identical to the nesting we draw. But we cannot have it both ways. We have cited this literature as independent corroboration when the direction of the result suited us, and the same citation obliges us to report that the corroboration is thinnest precisely where we most wanted it. More citation will not settle the matter. A coding scheme for organizational recursion would, if another team could apply it to the same cases and reach the same levels. Chapter 14 sets out what such an evidence base would have to contain, and Chapter 9 takes up the second finding Ostrom's data forces on us, about viability confirmed only in retrospect.

The one attempt to measure it used the organization chart

Somebody has now had to put recursion into an instrument, and what he did with it is the clearest evidence we have.

Michael Pfiffner's 2017 Utrecht dissertation operationalizes the Viable System Model into a questionnaire and applies it to 135 organizations. It is the most serious quantitative work this tradition has produced, and Chapter

[14]The principle is Ostrom's eighth, nested enterprises: Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), 90 (table 3.1). Michael Cox, Gwen Arnold, and Sergio Villamayor-Tomás, "A Review of Design Principles for Community-Based Natural Resource Management," *Ecology and Society* 15, no. 4 (2010): art. 38, https://doi.org/10.5751/ES-03704-150438; and, for the 2016 configural analysis, Jacopo A. Baggio, Allain J. Barnett, Irene Perez-Ibarra, Ute Brady, Elicia Ratajczyk, Nathan Rollins, Cathy Rubiños, et al., "Explaining Success and Failure in the Commons: The Configural Nature of Ostrom's Institutional Design Principles," *International Journal of the Commons* 10, no. 2 (2016): 417–39, https://doi.org/10.18352/ijc.634. The per-principle support ratios reported by the 2010 review are quoted inconsistently in the secondary literature, so we give them here from table 3 of the paper itself, as ratios of supportive to unsupportive cases: 4B, accountability of monitors, 11.7; 2B, proportionality, 10.0; 1A, user boundaries, 5.3; 2A, congruence with local conditions, 5.0; 4A, presence of monitors, 3.8; 6, conflict resolution, 3.2; 3, collective choice, 2.8; 5, graduated sanctions, 2.3; 1B, resource boundaries, and 7, right to organize, 2.2 each; and 8, nested enterprises, 2.0. Nesting is therefore alone at the bottom and not, as we had it in an earlier draft, tied with the right to organize. It is also the only one of the eleven for which Fisher's exact test failed at the five per cent level (P = 0.065; the authors report it as significant at the 10 per cent level), and it carries the smallest effect size (phi = 0.262). The authors' own bands treat ratios of 2 to under 5 as moderately well supported, so nesting sits at the very edge of the weakest band they use. In the 2016 analysis nested enterprises co-occur less with the other principles in successful cases (Baggio et al., 426, 428), and the successful irrigation and forestry cases include a solution set that contains every principle except nesting (429).

18 examines it at length.[15] The recursion principle is one of the properties he has to measure, and two items cover it. The first asks whether the system in focus's "overarching units (e.g. company, division, holding company) are themselves able to act as independent organizations at their level with all necessary functions which can fulfil their remit independently." The second asks the same of its "subordinate units (e.g. department, area, group)."

Beer's criterion survives in the item: the contained system must be viable in its own right. The discovery step does not survive. The organization chart supplies the levels to the respondent, in its own vocabulary: company, division, holding, department, area, group. The instrument does not ask where the levels are. It tells the respondent where to look, using the formal hierarchy as the answer, and then asks whether what is found there is viable.

We do not fault the design for this; nothing else was available. A questionnaire cannot ask an informant to derive a recursion level from first principles, because the model supplies no procedure for doing so, which is the charge of this chapter. Faced with the need to measure, the most careful attempt in the literature substituted the thing the model was supposed to improve on.

Beer had done the same thing first, and corrected it after the fact. His 1971 model for the Chilean state sector drew the recursions along lines of ownership. Events, he wrote later, "powerfully suggested that two levels of recursion were mixed together in that preliminary model," and a new level had to be inserted.[16] The first drawing followed the formal structure; the second followed what happened. Neither was produced by a ruler. Beer had warned against the first move in his own terms: if recursion becomes "a meaningless appendage to the term *organizational level*, then we do not have a valid model."[17] A recursion is not an organizational level, and the managerial systems at a level are not viable systems: "Systems Two, Three, Four, and Five are not in themselves viable systems at any level of recursion."[18] The neighbors in this chapter, Jaques's strata and Ostrom's nested enterprises among them, describe a related object, not the same one. Chapter 10 states his criterion once and compares each neighbor to it.

One number completes the point. Of thirty-five independent variables in that study, exactly one exceeded the threshold for missing data: the

[15] Michael Dominik Pfiffner, *System viability of organizations and the aetiology of organizational crisis: A Quantitative Assessment of Stafford Beer's Viable System Model* (PhD diss., Universiteit Utrecht, 2017), items P.1 and P.2, table 17, p. 165, and in the questionnaire reproduced in the appendix, p. 272; the missing-data figure is reported in the results chapter, §5.2.1, p. 180, where P.1 (UpRec) is identified as the only variable to exceed the five per cent threshold, at 5.9 per cent (eight values). **Read in full text 14 September 2026;** pages checked 23 September 2026 against the calibre copy of the Utrecht repository (publisher) version. Chapter 18 sets out the study's design and its limits.

[16] Beer, *Brain of the Firm*, 322–24.

[17] Beer, *Heart of Enterprise*, 308. On the organization chart see also Beer, *Diagnosing the System*, 2: it "makes no attempt to model VIABILITY."

[18] Beer, *Heart of Enterprise*, 411; see also 119, where the service departments "are not viable systems now, because their role is to serve the corporation."

upward recursion item, at 5.9 per cent. Every other variable came in below five per cent, most far below. Informed professionals with close knowledge of the organization, among them bankers, recovery officers and consultants, were asked whether the units above this one could stand as viable systems in their own right, and they left that question blank more often than any other in the instrument. The number is small and should not be made to carry more than it can. It suggests that the ladder is hardest to see in exactly the direction where the model requires it to be visible.

What this chapter hands forward

Koestler supplied the missing noun; Jaques, a ruler whose validity remains contested; Holling, phases that are difficult to identify in the present tense; PROSA, a graceful-degradation contrast. Beer supplied a criterion he called precise and rules of thumb that resolve to a judgment call. None of them supplied an instrument for discovering levels rather than drawing them. Chapter 11 draws the joint result: five traditions saw the structure, and none could measure it.

That absence is why the criterion, good as it is, cannot do the work we ask of it. A test tells us whether a proposed level survives scrutiny. It does not generate the proposal, and it does not adjudicate between two proposals that both survive. Chapter 3 finds the same pattern in our handling of requisite variety: an idea we treat as a law works in practice as a vocabulary for describing what we have already decided. In both cases the remedy is a cheaper experiment, not a better argument.

Receipt — blind comparison of a diagnosis and an interaction matrix. Take one organization with an existing VSM diagnosis and usable message, transaction or incident data. Register the block-detection rule before either analyst sees the other's result: a nested degree-corrected stochastic block model, with the number of levels chosen by minimum description length, which leaves no free threshold to tune after the fact.[19] One analyst receives the interaction matrix; the other receives the diagnostic material. Compare the detected blocks, the diagnosed recursion levels and, where records permit, Jaques's time-spans of discretion. Price: two analyst-weeks plus data preparation. If the blocks do not coincide, the negative result is that the diagnosed levels have no proxy in observed interaction structure. The full multi-organization design returns in Part Five, and Chapter 29 places this among the two cheapest experiments available to us.

[19]The method is to be confirmed with a network methodologist before registration. For the degree-corrected model, Brian Karrer and M. E. J. Newman, "Stochastic Blockmodels and Community Structure in Networks," *Physical Review E* 83, no. 1 (2011): 016107; for the nested version with the number of levels chosen by description length, Tiago P. Peixoto, "Hierarchical Block Structures and High-Resolution Model Selection in Large Networks," *Physical Review X* 4, no. 1 (2014): 011047. We have not read either paper in full.

Chapter 6

The Brain in the Vat

Philosophers of mind use the brain in the vat to name a cognitive system modeled without a body and without a world: everything the system needs to think is inside it, and the environment enters only as signal.[1] The phrase describes, with uncomfortable accuracy, the shape of a diagram we draw several times a week.

The diagram stops at the neck

We draw the model as a nervous system: channels, balances between them, a function that scans an environment, and an identity placed at the top of the page. The vocabulary is neurological and so is the picture, and both stop where a neck would be.[2] Beer's model does have a body: the body of the organization, whose organs are the operating units, kept stable by something like an autonomic nervous system.[3] The bodies of the people who staff it are below the line, and the picture does not reach them.

An organization, however, is enacted by bodies. It is carried by people whose knowledge is uneven and partly tacit, whose health varies, and whose time is finite and contested; people who fear specific colleagues and specific consequences, who have an appetite for some kinds of work and not others, and who have will. Everything the model describes actually runs on that substrate, so these facts are more than decoration on an implementation. Yet nowhere in the diagram does anybody want anything. Nothing corresponds cleanly to the capacity of the people expected to keep the system viable, whether physical, cognitive or emotional. As far as the notation is concerned, channels have capacity and the humans staffing them do not.

The claim needs to be stated precisely. It is a claim about the diagram and the practice by which we teach it, and Beer's books do not support a stronger one. He writes that the users of an operations room are "human

[1] Hilary Putnam, *Reason, Truth and History* (Cambridge University Press, 1981), chap. 1, "Brains in a Vat."

[2] Our description of our own teaching (Krishan's and Meg's course material); compare the diagrams in Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), and in Jon Walker, *The Viable Systems Model Guide*, 3rd ed., version 3.2 (2020), https://vsmg.lrc.org.uk/, cited in Chapter 1.

[3] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 75–76, on the heart and liver as the example of necessary autonomy, and 103 on the autonomic system that keeps "a stable internal environment."

beings constrained by their own neurophysiological limitations," whom it would be absurd to flood with data they cannot take in. He notes that divisional managers are "often unwilling" to learn the metasystem's language, and calls participation a matter of "the psychology of the system (and not its logic)." In his late account System Five supplies will and judgment, and the Chilean algedonic meter let a person signal a satisfaction "that does not have to be 'rationalized'."[4] Then, in the workbook written for diagnosis, he sets all of this aside on purpose. *Diagnosing the System* leaves out psychology, social anthropology and industrial relations, although "a lot is implicit in this book that derives from these fields."[5] Human capacity and will appear in the books in passing, and nowhere in the five-system notation or in the diagnostic method. That is the gap we inherited, and we taught it as if it were the model.

Meaning named, not modeled

The second form of the charge needs the most care, because Beer's texts say more about meaning than our teaching does.

Variety is a count of distinguishable states. That is what the concept is and what it is good for, and Chapter 3 sets out both the strength and the cost of taking it literally. Organizations, though, do not run on counts. They run on interpretation: on what a signal is taken to mean, by whom, against what history, and with what consequence for the person doing the taking. A figure that arrives in a monthly report is a state of the world already read by somebody with an interest in how it reads.

Three bodies of work locate the missing layer with some care. Bateson called it the difference that makes a difference, which is precisely the distinction between a difference that is merely available and one that is registered as mattering. Vickers located it in appreciative judgment, and argued against the whole regulatory framing that our tradition adopted. On his account, what an institution mostly does is maintain and revise the standards by which it reads its own situation, rather than correct deviations from a set point. Luhmann made meaning constitutive of social systems rather than incidental to them.[6] Whatever we make of any one of these, together they mark out territory that Beer names and does not model.

[4]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 243; Beer, *Brain of the Firm*, 159 and 283; Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150, for System Five as will and judgment; and *Heart of Enterprise*, 91, where "human beings become tired or distracted. At those moments, requisite variety is lost." The manager's "psychological space" and perceptual filter are in *Heart of Enterprise*, 381–93.

[5]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), ii.

[6]Gregory Bateson, *Steps to an Ecology of Mind* (Chandler Publishing, 1972), for the difference that makes a difference; Geoffrey Vickers, *The Art of Judgment: A Study of Policy Making* (Chapman & Hall, 1965), for appreciative judgment and the standards by which an institution reads its own situation; Niklas Luhmann, *Social Systems*, trans. John Bednarz Jr. with Dirk Baecker (Stanford University Press, 1995), chap. 2, "Meaning," and *Organization and Decision*, ed. Dirk Baecker, trans. Rhodes Barrett (Cambridge University Press, 2018), for meaning as constitutive of social systems.

He names it more often than we remember. "A system is not something given in nature, but something defined by intelligence." Whether a change counts as a state of the system depends on whether it "serves, or has no bearing on, the purposes of the system," and "everything depends on how the system is defined, and this in turn depends on who is the observer." Information "is what changes us"; data are its raw material, and "Data of themselves have no value." At each boundary "the 'language' of each subsystem is unique to it," so messages have to be translated. In *Beyond Dispute* (1994), "it is the interpretation of data that procures purpose."[7] Beer's sentence on information does the work of Bateson's difference that makes a difference.

What Beer states in principle, the notation does not carry. The five systems have channels and capacities and no place for interpretation, which is needed constantly, at every point where something is declared relevant, out of range or worth escalating. The diagnostic method starts from an act of interpretation, "Choose and clearly define the System-in-focus,"[8] and gives no account of how that choice, or any later reading, is made. This is the partition problem Chapter 3 raises, in a different guise. Distinctions do not come pre-labeled. Someone decides what counts as a state, what counts as a disturbance, and what a given reading means, and Beer says so. The diagram drawn from his model begins after that decision has been made and presents the result as structure. In our teaching, the point at which the observer entered drops out of view.

Two failures the book treats separately connect here. Chapter 2 argues that System Four is not a pair of eyes, partly because scanning presupposes a prior settlement about what is worth looking at. Chapter 25 asks what a transducer actually does when a signal crosses a boundary; on our reading, most of the work it is credited with is interpretive work we have not modeled.

Identity at the top compounds it

Placing identity at the apex of the diagram makes the meaning problem worse. It offers the analyst a resting place. Coordinated action can be explained by installing a smaller interpreter inside the organization, something that already knows what the whole is for and what its situation means, and whose knowing is then transmitted downward. This is the homunculus move. It explains nothing, because the interpreter's own interpreting is exactly what needed accounting for.[9]

[7] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 242 and 274; Beer, *Diagnosing the System*, 98; Stafford Beer, *Platform for Change* (John Wiley & Sons, 1975), 223; Beer, *Heart of Enterprise*, 101; Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 10.

[8] Beer, *Diagnosing the System*, 16.

[9] For the general form of the objection, Anthony Kenny, "The Homunculus Fallacy," in Marjorie Grene, ed., *Interpretations of Life and Mind: Essays around the Problem of Reduction* (Routledge & Kegan Paul, 1971). We have not yet read it; pages will follow when we have.

Beer's texts resist this reading more than our diagrams do. System Five "consists of many people," an "elaborately interactive assemblage of elements." Its main problem, he writes, "is its belief that it is *not* the puppet of some other system ... But it always is," because its input is mediated by Systems Three and Four, and "*the senior controllers inevitably work within the framework set by their juniors*." It sets the rules "not so much by stating them firmly, as by creating a corporate ethos," and Beer disowns the "connotations of seniority and autocracy" that drawing the metasystem on top carries.[10] Against this stands the language of *Brain of the Firm*, where System Five is "the thinking part ... the cortex itself" and a "thinking chamber" originates intentions.[11] On balance the objection lands on our practice and on that cortex language, and not on the System Five of *The Heart of Enterprise* and *Diagnosing the System*. What remains open is the specialist literature: we have not identified a statement of the homunculus objection against the VSM in the form we would need.

What would settle it

As it stands, the charge is a reading, and readings are cheap. Chapter 9 explains why our tradition accumulated so many of them: nothing we asserted was ever priced, so nothing was ever tested. This one can be priced.

Receipt — one comparison paper. The test is a systematic mapping of each of the five systems against the commitments of 4E cognition[12] — embodied, embedded, enacted, extended — stating for each component what changes under those commitments, what survives them intact, and what cannot be translated at all. Price: one focused review paper. The negative result must be allowed to count. If a faithful mapping leaves every component intact, the negative result defeats the disembodiment charge and the charge should be withdrawn.

We would regard that outcome as a good one. It would mean the model is more robust than we suspect, and that the resources for handling embodiment and meaning were present all along in a form we had failed to teach. Either result advances the position; only the absence of the test leaves us where we are. Chapter 28 sets out the general form of this commitment. The specific form here is simple: a criticism we are unwilling to lose is not a criticism. It is an allegiance.

Two limits remain. First, the mapping we are asking for is comparative, not confirmatory. A paper that sets out to show the VSM was 4E all along will not do the job, because the translation failures are the informative part. Second, this chapter has more unverified ground than most in Part

[10] Beer, *Heart of Enterprise*, 263–64 and 200; Beer, *Brain of the Firm*, 205 and 238; Beer, *Diagnosing the System*, 124.

[11] Beer, *Brain of the Firm*, 201 and 135.

[12] Albert Newen, Leon De Bruin, and Shaun Gallagher, eds., *The Oxford Handbook of 4E Cognition* (Oxford University Press, 2018), https://doi.org/10.1093/oxfordhb/9780198735410.001.0001.

One, and a reader is entitled to hold that against us. Chapter 7 argues that the model never became the machine it was described as being. The argument here is narrower and, for now, weaker: Beer wrote about bodies and readers, the notation he left us acquired neither, and we know exactly what it would cost to find out whether it can.

Chapter 7

A Machine That Never Became One

In 1963 Magoroh Maruyama pointed out that cybernetics had studied deviation-counteracting processes predominantly and paid less attention to deviation-amplifying ones, and he proposed a second cybernetics for the latter.[1] Beer's model was still forming when Maruyama wrote. Writing in 1979, Beer said that the neurocybernetic model "took over twenty years to develop."[2]

The VSM that came out of those twenty years describes the organization as a machine already assembled. Its five systems explain how an existing identity is maintained under disturbance: variety attenuated and amplified across an established boundary, operations coordinated without central congestion, and a purpose that converges between what is declared and what the system delivers.[3] They do not explain how a system becomes viable in the first place, how it passes through collapse as a phase rather than as a failure, or how it changes its kind. We have used the model in all three situations anyway, and we have rarely said what it cannot see there.

What the five systems are built to explain

Take the ordinary diagnostic occasion. An organization exists, its participants can name its identity, and it is handling some disturbance badly. The model gives us leverage because identity is fixed for the duration of the analysis. System Five's normative closure, System Three's resource bargain, System Two's damping and System One's autonomy within constraint are each defined relative to something the organization already is. Nothing

[1] Magoroh Maruyama, "The Second Cybernetics: Deviation-Amplifying Mutual Causal Processes," *American Scientist* 51, no. 2 (1963): 164–79.

[2] Stafford Beer, *The Heart of Enterprise* (John Wiley, 1979), xi. The next sentence reads: "During that time, but even more intensively since *Brain* was published, it has been used in every kind of organization in many different countries." The twenty years are not dated to end in 1972. Beer traces parts of the model to the "Sketch for a Cybernetic Factory" in *Cybernetics and Management* (English Universities Press, 1959), ch. XVI, and to the mathematical model of coupled homeostats he offered at a symposium in 1960 ("Towards the Cybernetic Factory," in *Principles of Self-Organization*, Pergamon, 1962): *Brain of the Firm*, 2nd ed. (John Wiley, 1981), 144, 264, 410.

[3] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 99: a comparator "continuously compares the declared purpose with the purpose imputed from the results that the system delivers," and the system "WILL CONVERGE ON A COMPROMISE PURPOSE." Maintaining the status quo as an end in itself is, for Beer, a pathology: Beer, *The Heart of Enterprise*, 277, 410.

about this is hidden; it is the plain content of the construction. The trouble lies in the range of cases to which we have applied it without qualification.

The model is homeostatic

Maruyama's date matters to how we tell the story, though nothing in the argument depends on it. Contemporaneous or not, the VSM as we teach and apply it is a first-cybernetics construction, though Beer's own texts are not first-order in this sense: they make the system an observer's act and closure a self-referential process.[4] It explains how a system holds together under perturbation.

It has no account of morphogenesis: of how a system becomes something it was not, of collapse as a phase rather than a failure, of the generative disorder from which new forms arrive. Panarchy has that account.[5] The adaptive cycle's release and reorganization phases have no counterpart among the five systems. No subsystem has the function of letting structure come apart, and nothing represents the interval in which a form is between kinds. System Four is where one might expect the capacity to sit. Beer gives it more than we teach: it contains environmental variety and also generates it, holds a model of the whole that includes itself, and is asked how it will "invent the future."[6] It is still not given the task of changing the system's kind. Beer's viable system undergoes "every kind of change. But there is no alteration."[7] Chapter 2, *System Four Is Not a Pair of Eyes*, establishes that boundary in detail, and we do not re-argue it here. Related asymmetries appear elsewhere in Part One. Chapter 4, *Channels Without Owners*, deals with revolt and remember and the couplings panarchy specifies and we never did; Chapter 5, *The Ladder Without a Ruler*, deals with temporal nesting. Across those chapters the pattern is consistent enough to name as one defect rather than three.

[4] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 242–47; Beer, *Heart of Enterprise*, 260, on "a *self-referential process*"; Beer, *Diagnosing the System*, 4 and 130.

[5] Lance H. Gunderson and C. S. Holling, eds., *Panarchy: Understanding Transformations in Human and Natural Systems* (Island Press, 2002); C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405. Chapter 4 sets out the cycle's four phases.

[6] Beer, *Heart of Enterprise*, 237; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 198, "a tool for *inventing* the future"; Beer, *Diagnosing the System*, 115–16, on self-awareness.

[7] Beer, *Heart of Enterprise*, 405; see also Stafford Beer, preface to Humberto R. Maturana and Francisco J. Varela, *Autopoiesis and Cognition* (D. Reidel, 1980), 70, where a complete change of oneself at any level of recursion is "not only difficult to accomplish but actually impossible."

An anatomy applied to things that are becoming

Process theory and panarchy handle dissolution and transformation better than the VSM handles either.[8] The sentence is easy to overread. An anatomy of maintenance is not defective for being an anatomy of maintenance. A model that explains how identity is held under disturbance has done something difficult and useful, and it is not answerable for questions it never set out to answer. Frameworks built for other phases do not weaken Beer's construction.

The defect is ours. We have taken an anatomy of the assembled organism and applied it, without qualification, to organizations that are forming, dissolving or transforming: startups before they have an identity to protect, institutions in the middle of dismantling themselves, bodies that came out of an intervention as a different kind of thing than they went in. In those engagements the model still produces a diagnosis. It reports missing channels, a weak System Two, System Three overreach. It does not report that the object under examination was not the kind of object the method addresses. We have not written that qualification into our practice, our teaching or our reports, and Chapter 9, *Why None of This Ever Cost Us Anything*, explains why the omission has been comfortable to leave in place.

Absorbing the vocabulary is not the same as making the connection

Beer took autopoiesis into *The Heart of Enterprise* himself and made it a diagnostic criterion. The time a system spends on its own self-production should be "exactly the time required to maintain System One," and Systems Two to Five "should exhibit no internal autopoietic behaviour whatever"; past that lies the pathology he calls the Establishment.[9] That is a use of the concept, more than a borrowing of the word. It concerns how much a system may spend producing itself, and says nothing about how a system becomes another kind of system. The wider field later took von Foerster and Maturana into its theoretical vocabulary.[10] Second-order formulations and the language of autopoiesis are now unremarkable in our discourse, and they do useful work in situating the tradition. A vocabulary, though, does not turn a diagnostic procedure that records none of the analyst's own choices into an account of autopoiesis,[11] and the presence of the terms cannot stand in for the demonstration.

---

[8] For process theory, Haridimos Tsoukas and Robert Chia, "On Organizational Becoming: Rethinking Organizational Change," *Organization Science* 13, no. 5 (2002): 567–82. We have not yet read it in full.

[9] Beer, *Heart of Enterprise*, 408, 411–12.

[10] Humberto R. Maturana and Francisco J. Varela, *Autopoiesis and Cognition: The Realization of the Living* (D. Reidel, 1980), https://doi.org/10.1007/978-94-009-8947-4, which carries a preface by Beer (63–72); Heinz von Foerster, *Understanding Understanding: Essays on Cybernetics and Cognition* (Springer, 2003), https://doi.org/10.1007/b97451.

[11] Beer, *Diagnosing the System*, 9 and 146, where the systems other than the next-level recursions are dedicated to stabilizing the internal environment and every homeostatic loop

A rigorous connection would have to say two things. The first is how the diagnostic observer takes part in producing the distinctions being diagnosed: how the consultant's choice of system-in-focus, recursion level and identity statement enters into what is then found. Chapter 6, *The Brain in the Vat*, sets out why that participation is not a marginal correction. The second is how a change of organizational kind appears in the method: what the diagnostician would see, in the model's own terms, when the object stops being one sort of system and becomes another. On the first, Beer went part of the way. He stated the observer's role in 1966, before the second-order vocabulary arrived, and his workbook compares members' perceptions: check the First Axiom "as System Three sees it; as System One sees it … There are ALWAYS surprises."[12] He gives no procedure that exposes the consultant's own choices of focus and level to a second reader. Neither that procedure nor a representation of kind-change has been shown in the material we have examined here. Nor have we run and reported a dated search of the literature for work that might already have shown it. A positive result would need more than a paper citing Maturana alongside Beer. It would need a procedure that makes the observer's contribution explicit, or a representation of kind-change in the five systems, applied to a case and legible to a reader who did not write it.

The consequence, stated exactly

For any diagnosis of an organization that is forming, dissolving or transforming, the limitation is exposed: a model of survival cannot, by itself, say what is worth becoming. The claim is narrower than it may sound. It does not say that viability is the wrong thing to care about, and it is not an argument about values; Part Four takes up the related question of what survival is *for*, and Chapter 23, *The Narrow Claim*, sets the boundary of what we are prepared to assert. The charge here concerns representation. An organization can be in states the model has no way to depict, and we have issued diagnoses in those states without saying so. Chapter 27, *What Might Replace the VSM*, considers what a client is entitled to expect in them.

Beer's own record contains one such state. After October 1972 in Chile, he wrote, "there seems to be a resurgence in the self-organizing capability of communities to promote themselves as viable systems." He handled it by redrawing the recursions after the event, with a new level and three new Systems One.[13] The method caught up with the change; it did not represent

is to be checked against the four principles. Beer qualifies the method himself: 6, a system-in-focus may lie on several recursive dimensions at once; 99, "A GOOD OBSERVER will impute the purpose of the system"; 144, "There is no 'correct' interpretation of the VSM."

[12] Beer, *Decision and Control*, 242–47; Beer, *Diagnosing the System*, 95. See also Beer, preface to *Autopoiesis and Cognition*, 67: "A system is not something presented to the observer, it is something recognized by him"; and Beer, *Heart of Enterprise*, 36: "The observer is always part of the system: he determines its nature, its purpose, its variety."

[13] Beer, *Brain of the Firm*, 322–26; the quotation is at 323.

the change while it was happening. Beer also put the identity question himself: after adaptation, will the system "still be able to recognize itself?"[14]

Receipt — specify kind-change. The way to settle this is to attempt the specification and see what it costs. Write one paper stating what System Four would have to observe, represent and alter in order to anticipate a change in the system's kind rather than a change in its environment. Price: one conceptual paper tested against at least one documented transformation, beginning with Beer's own redrawing of the Chilean recursions after October 1972. If no operational distinction can be made, the negative result is that the extension adds vocabulary but no diagnostic power — which is worth knowing, and worth publishing, because it would tell us that the limitation is structural rather than an unfinished piece of work. Either outcome ends the present situation, in which we neither have the capacity nor concede that we lack it. The Coda, *Judy*, returns to what we owe the organizations that were becoming something while we were describing them as though they already were.

[14]Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 149–50.

Chapter 8

The Loop That Stayed an Intention

Project Cybersyn promised a line running from worker participation to government action. Beer's own account documents that line as an intention and a procedure. The channels actually built ran one way, which is why the missing loop looks designed out rather than accidentally omitted.

The promise and the documented design

Chile is our founding case, and we use it for something specific. When the Viable System Model is asked to show that it is more than a management diagram, that it belongs to a politics as well as to a control theory, we point to Santiago: a national economy, a government that asked for the work, a designer who described what he was building as an instrument of the workers' own self-management. The case carries a great deal of argumentative weight for us. Most of us have never gone back to check what the design specified, or what Beer himself reported about how it ran.

The channels that were specified run in one direction. Indices are drawn from the enterprises and filtered as they rise. Forecasts are computed centrally, although Beer's design promised each level its own tools and its own computer.[1] Exceptions are raised to a level above the one that produced them, and the government, at the top of the structure, can act on what arrives.[2] Each of those channels is described with care. The channel that would run the other way is described differently, as a procedure and an intention.[3] The teams modeling each plant were to build the model with the people who worked there and agree the performance measures with them, and Beer records that "this was not always being done in the

[1] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 275: each unit receives "the whole *Cybersyn* package: the four tools, and a computer," and "any given level of recursion receives directly only its own Cyberstride reports." Whether this held in practice is the question Medina and Shalizi address.

[2] Beer, *Brain of the Firm*, chap. 17, "Into its Stride," 259–310, especially 275–76 on Cyberstride, the recursions and algedonic escalation; Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 71–74.

[3] Beer, *Brain of the Firm*, 252: "the intention of Cybernet was to make computer power available to the workers' committees in every factory"; 271 on modeling and measures agreed with the people in each plant; 276, where the models and the recovery times were to be made "with the help and the agreement of workers' committees"; 278–84 on the People Project and the algedonic meter; 324 for the redesigned recursion under "the workers' committees."

intended spirit."[4] The algedonic meter, the one channel from the people to the government, had a wall reserved for it in the operations room: "there remained a blank wall in the Operations Room: it was intended to house an algedonic meter." We read Beer's sentence as saying that the meter was never installed.[5] What never existed is a route with standing, by which the workers the design names as its client could contest an index, revise the model of their own enterprise against the modelers' view, or limit what the government does with the reports.

We cannot treat that as an omission awaiting a later draft. Everything adjacent to it was built. A design in which the client's return channel is the one element left as a procedure and an intention tells us what it takes the client to be: a source of measurements and a partner in modeling, without a part in the decision. That is the missing loop, and it is our problem before it is anyone else's, because we have taught the case for fifty years as though the intention had been the channel.

Twenty-seven pages, and a tape

In 1981 Werner Ulrich publishes twenty-seven pages in the *Journal of Applied Systems Analysis*. They put to Project Cybersyn five questions C. West Churchman had formulated for any systems study: who is the client, what is the purpose, who is the decision maker, who is the expert, and who is the guarantor.

He is careful about what he is not doing. "It is not my intent," he writes, "to question the social and democratic ideals that both Salvador Allende and Stafford Beer associated with the project." He revises the paper after comments from Fernando Flores, Allende's Secretary of Economics. An anonymous reviewer who worked on the project confirms its accuracy.

His conclusion is one sentence: "Cybersyn's alleged goal of devolving power to workers is not confirmed by its design."

Beer sent his comments on tape. He could not agree with the critique. Nothing from the tape was published, and the tape, if it survives, has not been found in the catalogue of Beer's papers. Two years later Beer published a short reply in the same journal, which we have not yet read.[6]

Apart from Beer's own five pages, that is the whole of this field's recorded reply to Ulrich's structural argument.[7]

---

[4] Beer, *Brain of the Firm*, 271.

[5] Beer, *Brain of the Firm*, 283–84; the blank-wall sentence is on 284. Beer describes a prototype of ten linked meters built in England and taken to Chile for experiments (284), not an installation in the operations room.

[6] Stafford Beer, "A Reply to Ulrich's 'Critique of Pure Cybernetic Reason'," *Journal of Applied Systems Analysis* 10 (1983): 115–19. We have not read it. It is listed and quoted in M. C. Jackson, "An Appreciation of Stafford Beer's 'Viable System' Viewpoint on Managerial Practice," *Journal of Management Studies* 25, no. 6 (1988): 557–73, at 572 and 564.

[7] Werner Ulrich, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59 — twenty-seven pages. All the quotations, the revision after Flores's comments, the confirmation by a reviewer who

Beer did answer, first in the medium he chose, as the acknowledgements record, and then briefly in print. What is missing is anything the field can read. A critique built with the cooperation of the people who ran the project, published in a systems journal, addressed to the founder of the model and to the case we cite most often, has gone forty-five years without a reply in print from any of us. The omission belongs to us, not to Beer.

Beer did engage the political content of the charge in print, in the same year and without naming Ulrich. The second edition of *Brain of the Firm* concedes that the lapse from participatory modeling "seemed likely" to have "aided detractors who later wished to call Cybersyn 'technocratic'," describes a current inside the project that was "technocratically minded," and calls some of the hopes for handing the tools to workers' groups "manifestly starry-eyed."[8] That is a concession. It does not answer Ulrich's argument about what the channels were wired to do.

What the critique actually says

Ulrich did not write from a distance. "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics" is a case study of Project Cybersyn built on Beer's own reports, on Fernando Flores's account, and on an insider reviewer who confirmed its accuracy.[9] The finding is not that

worked on the project, and the taped reply are from the paper and its acknowledgements; the note below on Ulrich sets out the argument in full and what its standing rests on. The five categories are C. West Churchman's, applied by Ulrich. The search for the tape so far: Liverpool John Moores University, catalogue of the Stafford Beer Collection (November 1996), searched on 29 September 2026 for "Ulrich", "critique of pure", "Applied Systems Analysis", "tape" and "cassette", with no entry; the collection itself and Ulrich's papers have not been searched.

[8] Beer, *Brain of the Firm*, 271, 322, 328.

[9] Werner Ulrich, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59. A Spanish translation appeared as *Una crítica de la razón cibernética pura (el caso del proyecto CYBERSYN)*, Memos de Investigación 36 (Universidad de los Andes, 1989). Earlier drafts of this book recorded that we had not been able to obtain the text and were relying on Ulrich's own one-line description of it; we have read it since, and this note replaces that admission. The paper has two halves. The first argues that systems theory has been absorbed by a "scientistic" school which mistakes theoretical for practical reason, and proposes a critically-normative paradigm against what Ulrich names the *syntactic fallacy* — taking Ashby's law as a criterion of systems rationality when the level of organization is socio-cultural rather than organic. The second applies five of Churchman's categories — client, purpose, decision maker, expert, guarantor — to Cybersyn, and concludes that "Cybersyn's alleged goal of devolving power to workers is not confirmed by its design. Instead, its strongly hierarchical organisation and its concept of 'autonomy' one-sidedly serve the top decision maker, the government." Three things should be said about its standing. Ulrich states at the outset that the political purpose Cybersyn served is *not* the subject of his critique — "It is not my intent to question the social and democratic ideals that both Salvador Allende and Stafford Beer associated with the project" — and closes by regretting that his argument falls on "one of the very rare efforts undertaken by systems scientists in the service of democratic values." The account was revised after comments from Fernando Flores, Allende's Secretary of Economics, and its accuracy was confirmed by "an anonymous reviewer who was involved in the project"; where Ulrich contradicts the contemporary *New Scientist* coverage, he does so on Flores's authority. And on the question of

Cybersyn failed. It is that the design does not do what its designer said it did. The client Beer declared was the workers. The client the design serves is the government, which alone can act on the system's reports without a homeostat above it, change the measure of performance, and control what goes in. Beer's own recursion supports the reading: it names the government as the viable system in focus, "System Five being the President of the Republic." His principle that "System Five within a people's government cannot be an élite ruling class: it is somehow the embodiment of the mass of the people themselves" names no mechanism.[10] Autonomy in that architecture holds only while things run smoothly, because an algedonic↗ signal is a transfer of control upward as well as a cry of pain. Hence the sentence in the paper that is hardest to answer, that "what distinguishes Cybersyn from a traditional state bureaucracy is not a less hierarchical structure but the greater speed at which vertical communication (and control) takes place." Part Four is built to answer that objection, made in 1981 and aimed at the founder.

Two things follow for the rest of this book. The objection is structural rather than historical: it would stand even if the project had run for a decade, because it concerns what the channels are wired to do. And it lands on the part of the model we have specified least. Chapter 4, *Channels Without Owners*, sets out how little the algedonic channel has been given, beyond Beer's statistical thresholds, in the way of owners, addressees or rights of reply; Ulrich's paper is the earliest place we know of where that gap is treated as a political fact rather than an engineering detail. Chapter 16 states his argument in full, together with Hayek's and Cooper's.

The strongest defense available

The best defense is that Cybersyn was never built to compute a plan. It attenuated. It filtered indices, forecast statistically, reported exceptions, and was meant to keep an alert with the unit for a defined interval so that the

a reply: the acknowledgements record that "Stafford Beer too was kind enough to transmit his oral comments on tape to me, though he could not agree with my critique." Beer replied in print two years later, in the same journal (10 [1983]: 115–19), which we have not yet read. His two 1973 letters in *New Scientist*, of 22 February (answering Joseph Hanlon) and 22 March (answering Herbert Grosch), answer that year's critics, not Ulrich (Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 182, 186–87, and 287 n. 40, 289 n. 55). The tape, if it survives, would be worth finding. One limit of the paper is worth stating in its favor as well as against it: it is built on *Decision and Control*, *Brain of the Firm*, *Platform for Change*, *Fanfare for Effective Freedom* (the Richard Goodman Memorial Lecture, Brighton Polytechnic, 1973) and "Cybernetics of National Development Evolved from Work in Chile" (the Zaheer Foundation Lecture, New Delhi, 1974), and engages neither *The Heart of Enterprise* nor *Designing Freedom*.

[10] Beer, *Brain of the Firm*, 249 and 288; for Allende's "*el pueblo*," 258.

unit could absorb its own disturbance.[11] That is a Hayekian architecture:[12] leave the decision to the person on the spot and escalate only the residue. Beer called his public lectures *Designing Freedom*, and meant it.[13]

We take that defense seriously, and quoting the phrase "cybernetic management" back at it does not dispose of it. On this reading the design is a subsidiarity machine. The interval before escalation is a right to be left alone for a stated period, and the filters protect the enterprise from the center's attention instead of delivering the enterprise to it. Chapter 26 takes up what Beer meant by freedom and whether it can be stated as a quantity. If the reading is right, the missing loop does less damage than it appears to, because the architecture would concede to each unit the decisions that matter most to it.

What the record cannot settle

The historical record neither confirms nor refutes the defense, because the system never ran at a scale that could have tested it. The most damaging criticism is empirical. Cosma Shalizi's verdict is severe, and we accept its force where the evidence supports it. To describe Cybersyn as a decentralized planning system "is nonsense", since it was computationally centralized on a single mainframe. Time lags made its alerting close to useless for early detection. Its safeguards for enterprise autonomy could, on Medina's account, have been undone as easily as by someone from the telex room walking down the hall. The only part that demonstrably worked was the telex network, which proved its worth coordinating trucks during the October 1972 strike, as logistics rather than as cybernetic management. At best it would have been a dashboard.[14]

[11] Beer, *Brain of the Firm*, 276: "the alerting signal of incipient change, which should stay with the autonomous Beta, will automatically convert into an algedonic signal … unless Beta succeeds in overcoming the problem," within "physiological recovery times for each index" that the operational research teams were to nominate. Medina, *Cybernetic Revolutionaries*, 73, describes the practice differently: the alert went to CORFO and to the factory's interventor together, and the government then gave the interventor "a limited window of time to resolve the problem on his own" before CORFO intervened; see also 184.

[12] F. A. Hayek, "The Use of Knowledge in Society," *American Economic Review* 35, no. 4 (1945): 519–30.

[13] Stafford Beer, *Designing Freedom* (CBC Learning Systems; John Wiley & Sons, 1974), delivered as the 1973 CBC Massey Lectures.

[14] Cosma Shalizi, "Project Cybersyn," Notebooks, http://bactra.org/notebooks/cybersyn.html, consulted 29 September 2026 (the page shows its last update as 1 August 2026); it is revised continuously. The quoted verdict is that "to describe it as a decentralized planning system is nonsense"; the supporting points are that there was "the one (!) mainframe computer available to the project," that the lag between measurement, processing and delivery "was so long that it seems to have been a complete flop," that "dismantling one of the primary safeguards of [enterprise] autonomy might have been as easy as having someone from the telex room walk down the hall," and that what the project would have produced at best was a "pioneering example of what we'd now call a 'dashboard.'" The "walk down the hall" sentence is Medina's own: Medina, *Cybernetic Revolutionaries*, 184; on the delays, ibid., 216. For the October 1972 strike, ibid., 141–69 (the telex network, 149–50); for the scale at which the system actually ran, ibid., 128 and 196.

Beer's account says otherwise on both points. He describes the strike as a test of filtration, in which "some signals were algedonic, requiring instant decisions, while others could be attenuated," and he puts the coverage at "something like seventy percent of the socio-industrial economy," about four hundred enterprises.[15] These are the founder's claims, made in an account he wrote in the first person and as a partisan. Medina's figures for the scale at which the system ran differ.

What that licenses is limited. It does not show that attenuation is the wrong architecture, or that the Hayekian reading of the design is mistaken. It shows that the founding case supplies no evidence for either, and that the safeguard on which the whole defense rests, the interval Beer's design gave a unit before CORFO intervened, which in practice ran while the alert had already reached CORFO, was on the available record a convention rather than a constraint. If a staff member could remove a protection by walking down a corridor, the design does not enforce it. Ulrich reached the same finding from the other side: the client's position in the architecture depends on the goodwill of the party above it.

We have kept the case in our teaching for fifty years without either finding attached to it. Chapter 9 asks why that was so easy, and its answer does not flatter us: nothing in our practice made the omission expensive.

The boundary condition

Ostrom's seventh principle supplies the boundary condition.[16] No amount of internal viability survives the withdrawal of the right to self-organize. That right is a precondition, prior to any design choice. The model places external authority in the containing recursion and in its command and legal channels, and Beer knew that such authority could dispose of a unit: his example is the division chairman "secretly moving to *sell* his division."[17] The model does not treat the revocation of that embedding as a variable with its own regulation. September 1973 does not test the VSM's performance; it demonstrates that an externally revocable right to organize decides whether any internal design gets to operate. The Shalizi material, likewise, bears on what the founding case documents and not on a performance test the project never reached.

We state this with the restraint the events require. Nothing in the model is answerable for what happened in Chile, and we make no claim about it here beyond the structural one: a system whose right to exist is held by

[15] Beer, *Brain of the Firm*, 312–14 on October 1972, and 264 for the coverage (at 311, "Roughly 60% of all enterprises"). The account is declared first-person at 245–46. For Medina's figures, see the note on Shalizi above.

[16] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), 90 (table 3.1), principle 7, minimal recognition of rights to organize.

[17] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 351; Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 1 and 38–41; Beer, *Brain of the Firm*, 275, on the metasystem's "authority to delimit the autonomy of any … System One."

another party has a variable outside its own regulation, and our diagrams do not represent it. Chapter 14 sets out what recognizing external authority as a design variable would require of us, and what evidence we would have to gather to say anything about it beyond assertion.

What would settle it

Receipt — find the tape. Read Beer's 1983 reply first. Then search Stafford Beer's papers at Liverpool John Moores University and Werner Ulrich's papers for the oral reply, recording catalog scope, search terms and date. Price: archive fees and one researcher-week. If the tape cannot be found, publish the reply the field owed Ulrich: reconstruct Beer's strongest answer from the contemporaneous record and submit it with Ulrich's unanswered claims stated in full. Forty-five years after 1981, the silence is itself part of the evidence, and it is ours to account for. We return to what it has cost us in the Coda.

Chapter 9

Why None of This Ever Cost Us Anything

The claim

The mapping we teach is not the mapping Beer wrote (Chapter 1). System Four is taught as a pair of eyes rather than as the function he specified (Chapter 2). The law we quote travels under a changed verb, and Ashby did not write it (Chapter 3). The channels are specified as mechanisms and never as institutions: the filter that decides what pain reaches the top belongs to the metasystem, and nobody has measured a channel (Chapter 4). The ladder of recursion has a criterion Beer called precise and no ruler, so two competent practitioners can count different numbers of levels in the same organization and neither can be shown to be wrong (Chapter 5). The notation has no place for the people who staff it, or for the interpretation Beer names in principle (Chapter 6). There is no account of becoming; the machine is analyzed as though it had always existed (Chapter 7). And the founding feedback loop was documented as an intention and never built as a channel with standing (Chapter 8).

None of these eight defects has ever had to count against the Viable System Model, because nothing in the model's empirical core, the necessity of the five functions and the recursion claim, has been stated in a form that could be disconfirmed, and we have never said what observation would make the model give something up.[1]

They are different kinds of error. Some are misattributions, some are omissions, one is an unfinished piece of engineering history. What they share is their fate. Each has been visible in the literature for years, some for decades, and none has cost the model a claim, a diagram, a teaching practice or a client engagement. Nobody examined them and found them

[1]The qualification is deliberate, because Beer did state some things in a form that invites a test. He set out the Recursive System Theorem and wrote that "the proof offered later will be a topological proof" (Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 118); we have not found that proof in the book. He claimed that the model gives "the necessary and sufficient conditions of any viable system" (ibid., 235), which is the claim a counterexample would have to hit, and the escape routes set out below protect it. He measured a model's usefulness by "straightforward testing of the predictive value of the scientific model" (Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113, paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994)), and the tradition kept his sentence about usefulness and dropped the test. And he expected that "maybe (one day) thresholds will be established, to exceed which will predict collapse" (Beer, *Heart of Enterprise*, 432); we know of none. None of these supplies an observation that would count against the core.

harmless; our literature has no procedure by which a defect could be found harmful. We have built a body of work that cannot be charged, and then treated the absence of charges as a good record.

Unfalsifiable in practice, though not in principle

Any organizational failure can be redescribed afterward as a missing system, a pathological system or an attenuated channel. The redescription is always available, because the vocabulary is complete. A framework that absorbs every outcome is a language rather than a theory. Languages are useful, and we need not be embarrassed to have built one, but a language does not accumulate the credit a theory accumulates, and evidence cannot improve it.

The repair is available and we have not made it: state what would disconfirm. The model's own claims already imply two candidates. A system that persists in a changing environment while demonstrably lacking one of the five would do it, because the model asserts that all five are necessary. So would a system that dies of a structural cause the model cannot see, with all five present, intact and healthy, because the model asserts that the five suffice to account for structural viability.[2] Neither observation is exotic. A research program of ordinary size could look for both in the field, and Chapter 28 sets out what such a specification would have to contain.

Until we look, the record is not clean. It would once have been fair to call it untested, but that is no longer accurate. The model has been tested, most recently by Schwaninger and Scheef in 2016[3] and by Pfiffner in 2017, and Chapter 18 examines both; two earlier St. Gallen dissertations, which we have not read, made the first attempts. The accurate charge is narrower and worse: in the two tests we have read, the model was put at risk neither time. In the first study the one hypothesis that failed was explained away after the result, against no rule stated beforehand. In the second the necessity claim sat where the statistical model could not reach it, so no possible pattern of results could have shown a subsystem to be dispensable.

A record produced that way carries no information about whether the model is right. It is compatible with a model that is right, with a model that is wrong, and with a model that says nothing determinate enough to be either.

---

[2] Beer, *Heart of Enterprise*, 408: "if a system has all of the necessary and sufficient conditions of viability, which we set out to discover and have now listed, it will live."

[3] Markus Schwaninger and Christine Scheef, "A Test of the Viable System Model: Theoretical Claim vs. Empirical Evidence," *Cybernetics and Systems* 47, no. 7 (2016): 544–69, https://doi.org/10.1080/01969722.2016.1209375. Page references are to the online-first version (published 12 August 2016 and paginated from 1), the copy we read. Hypothesis 3, on System Three, is rejected at 14; at 17 the authors "surmise" that the negative result does not reflect a loss of importance of the function and that it "has evolved new structural solutions." Chapter 18 examines the study.

The mechanism can be caught operating on itself, which is why this belongs here rather than in Chapter 18. Pfiffner opens by adopting critical rationalism explicitly: following Popper, he writes, all knowledge is hypothetical and "cannot be verified but only falsified." In his discussion chapter, among the merits of the model he is using, he lists that the Viable System Model "has a long track record of application and (peer) review and, to date, has not been falsified", and earlier he quotes Schwaninger to the same effect as part of the project's warrant.[4] A researcher who took the trouble to declare falsificationism cites, five chapters later, an absence of falsification as a reason for confidence. He is not being careless. The sentence is simply there in the literature, available to anyone who reaches for a reason to trust the model, and nothing in our field flags what it is.

Beer built the escape routes in

Both candidates meet an obstacle before anyone goes into the field, and the obstacle is in Beer's own books. The model claims to state "the necessary and sufficient conditions of any viable system."[5] The same books say that the functions a disconfirming case would have to lack are never really missing. System Two, the anti-oscillatory function, "is always present, or the organization would shake itself to pieces."[6] Most of the successful System Twos Beer had observed "have been entirely unofficial and largely unrecognized."[7] Of System Four he wrote that it "exists — it has to: but it is disseminated, and lacks integration," and that "there is always a System Four, even if it is not identified in quite the form specified here."[8] An organization that seems to lack either function has, on this account, an unofficial or a scattered one.

The second route covers survival. An institution whose metasystem has collapsed into System Three is, in Beer's image, a decerebrate cat: anaesthetized, the cerebrum removed, pinned to the table and fed. "It lives on; its viability is ensured by a bogus environment, and is sustained

[4] Pfiffner, *System viability of organizations*: the critical-rationalist position is stated in the introductory chapter, §1.4.2, p. 8, citing Popper and Keuth, *Logik der Forschung*, 11th ed. (Tübingen: Mohr Siebeck, 2005); the "has not been falsified" claim appears in the discussion of findings in the author's own voice, chapter 6, p. 222, and earlier in the theoretical framework, §3.4, p. 85, quoting Schwaninger (2000) in the author's own translation ("not been falsified until today"). **Read in full text 14 September 2026;** pages checked 23 September 2026 against the calibre copy of the Utrecht repository (publisher) version. Chapter 18 sets out the study in full.

[5] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 235.

[6] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 66.

[7] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 175, where the sentence follows the remark that "our age has produced so many large-scale organizations without a System Two."

[8] Beer, *Heart of Enterprise*, 143; Beer, *Brain of the Firm*, 183, which continues: "It will always be found as that set of activities, maybe disseminated, which feeds the highest level of decision making."

by artificial sustenance." Beer added: "And this is called 'living'."[9] He also refused to treat existence as proof. "The fact that it is there does not prove that it is effectively there, witness universities, nor efficiently there, witness hospitals."[10] So a system that persists while lacking a function can be read either as one that never lacked it or as one whose persistence is not viability. Each reading is reasonable on its own. Together they close off the first candidate unless someone decides in advance what counts as absent.

Beer stated his diagnostic method in the same form. Breakdowns "may be diagnosed, simply in the fact that some inadequacy in the system can be traced to malfunction in one of the five subsystems." His confidence in the model, "subjectively speaking," came "not so much from the fact that the pathology of the viable system can be investigated with ease, as from the speed with which the diagnosis can be made."[11] That is the redescription described above, in the founder's words and on the founder's warrant.

None of these moves is illegitimate. Functions do migrate into informal roles, and life on artificial sustenance is not health. What they require is a test that closes both routes before the data are seen: each function coded, by indicators fixed in advance, as official, disseminated or absent, and an outcome that is Beer's viability, the capability for a separate existence, rather than survival. That is the coding in Chapter 28.

The pathology catalogue is the same property seen from the other side

The same point can be made from the model's strongest ground. The VSM's pathology catalog is generated rather than collected. System 3–4 imbalance, the autonomy–cohesion oscillation, an algedonic↗ channel that is present but unheeded, identity that has quietly relocated into operations: all of these are derived from the structure, not gathered from cases. That is a real and unusual advantage. Most rival frameworks have case libraries where we have a theory, and a theory that predicts failure modes before anyone has seen them is worth more than a taxonomy assembled after the fact.

The concession belongs in the same breath. Generativity without testing is also the recipe for a vocabulary that cannot be wrong. The structural derivation that lets us name a pathology in advance lets us name one after any outcome whatever, and nothing in our practice distinguishes the two uses. The catalog and the defect above are one property, seen from two sides.

[9] Beer, *Heart of Enterprise*, 265, following "Then the whole metasystem collapses into System Three."

[10] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 148.

[11] Beer, *Think Before You Think*, 147–48.

Beer's own epistemology invites the question

Beer's epistemology was more modest than his manner. Models, he held, are built by observers for purposes, and the purposes belong to somebody; his own he submitted, in his phrase, cheerfully to Popperian criteria of disproof.[12] He also wrote the sentence the tradition quotes most: "a model is neither true nor false: it is more or less useful." And he attached a test to it. A model's utility is measured by "straightforward testing of the predictive value of the scientific model. If this predictive value is low, we had better start again."[13] We kept the first sentence and dropped the test. For sufficiency he had one test, to ask users whether anything necessary had been left out; the answer was "a nil reaction," which he called "not a logically rigorous test; it is merely an empirical verification."[14] On his account, asking why none of these defects cost the model anything is the question his philosophy invites, and putting it is a form of respect. A community that declines to ask it departs from Beer's position while believing it protects it.

Viability is confirmed only in retrospect

A second structure is at work, and it is not peculiar to us. If a set of principles can diagnose the long-lasting and the already-dead but not the young and viable, it describes what survival looks like in retrospect and is no instrument for the present tense. The finding arises in Ostrom's evidence base, where the design principles were drawn from institutions selected for having endured; Chapter 14 takes it up in full.

The VSM, as we practice it, makes the same move. Beer defined viability as a capability, being "able to maintain a separate existence," and not as the fact of survival.[15] He also held that the present tense could be read: "there are always pre-symptoms" of instability, and "nobody is measuring

[12] Beer, *Brain of the Firm*, chapter 20, "Prospectus," section "The Cybernetics of Crisis," p. 377, where he says of his models that "they are cheerfully submitted to Popperian criteria of disproof." A correction to an earlier version of this note, which said that "a model is neither true nor false" could not be found in Beer. It is his: "So remember: / a model is neither true nor false: / it is more or less useful," set as three broken lines in Beer, *Diagnosing the System*, 2. Variants: ibid., 144 ("more or less useful interpretations" of the VSM); Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 96, where a methodology is "neither right nor wrong, but only more or less useful"; Beer, *Think Before You Think*, 254: "A system is a model that is more or less useful for some purpose." The two positions point the same way, because Beer tied usefulness to prediction (see the next note).

[13] Beer, *Diagnosing the System*, 2, for the first sentence; Beer, *Decision and Control*, 113, for the test. On the same page Beer says that a model which ignores relevant elements "will lose in utility, but it cannot lose in validity."

[14] Beer, *Heart of Enterprise*, 115.

[15] Beer, *Diagnosing the System*, 1, where Beer takes the definition from the *Oxford English Dictionary*; see also Beer, *Think Before You Think*, 107: "A viable system is one sustaining the capability for independent existence as a recognizable identity."

stability. It can however be done."[16] The thresholds that would turn such readings into a forecast he left for later: "maybe (one day) thresholds will be established, to exceed which will predict collapse."[17] We know of no validated instrument of either kind. In practice, then, success is defined as viability, and viability is confirmed afterward. An organization alive today may be about to fail, and the model as we use it cannot say so in any way that does not depend on the outcome it is supposed to predict. That is why the redescription always works. Survivorship does not make a framework false. It makes its confirmations uninformative, because the cases were selected by the very property the framework claims to explain.

A contemporary parallel, offered without malice

We are not alone in this position, and naming a neighbor in it is no accusation. Friston himself, asked directly, has said of the free energy principle: "Like Hamilton's Principle of Stationary Action, it cannot be falsified. It cannot be disproven."[18]

The concession is honest, made in the author's own words and volunteered rather than extracted. It is also precise about what remains. What survives is a generative machinery for building testable process theories, and those process theories can be confronted with data in the ordinary way. What does not survive is the claim to be a law about the world that observation could confirm or refute. We recognize the position, because it is our own. The difference between the free energy principle and the VSM on this point is not the epistemic status of the framework. Beer came close to Friston's concession for the law at the model's core. When a friend called Ashby's law "a mere tautology," Beer told him to suspect the word *mere*, since "the entire corpus of mathematics is either tautologous—or wrong."[19] For the claim that the five systems are necessary and sufficient we have found no such statement. One community has said in print what its principle is. The other has said it only for one law.

---

[16] Beer, *Heart of Enterprise*, 374 ("there are always pre-symptoms of the fact") and 275 ("nobody is measuring stability. It can however be done. The key to the problem is statistical filtration").

[17] Beer, *Heart of Enterprise*, 432, where the exercise is meant to show "how successfully the law of cohesiveness is working in a given enterprise."

[18] Karl Friston, Martin Fortier, and Daniel A. Friedman, "Of Woodlice and Men: A Bayesian Account of Cognition, Life and Consciousness; An Interview with Karl Friston," *ALIUS Bulletin* 2 (2018): 17–43, at 21. The context is his distinction between a normative principle and a process theory: "the free energy principle is what it is—a principle. Like Hamilton's Principle of Stationary Action, it cannot be falsified. It cannot be disproven. In fact, there's not much you can do with it, unless you ask whether measurable systems conform to the principle." He goes on to argue that demanding falsifiability of a principle is a category error.

[19] Beer, *Think Before You Think*, 45; compare Beer, *Decision and Control*, 280: "if the law of requisite variety is a truism, how is it that we try to disobey it all the time?"

How the absence of objections is read

This has a consequence for how our work is received. A literature that contains almost no published results against itself is, in our judgment, read as a school rather than as a science, and, we suspect, cited less. Readers outside the field cannot tell a framework that has withstood testing from one that has not been tested, except by looking for the tests. When they look and find none, the reasonable inference is the unflattering one, and we have supplied it since *Brain of the Firm* appeared in 1972. Chapter 16 sets out the objections that do land, and they are fewer and narrower than the reception would suggest.

The count itself is unverified

Even our account of how thin the critical literature is remains unverified. The published critical canon may be as thin as two normative lines: Ulrich's charge of technocratic hierarchy and Amey's charge of excessive generality.[20] That count is not ready to carry prose as fact. Until it has been reproduced, "two" is a research proposition, not a finding.

Receipt — make the canon count reproducible. Search named scholarly databases for published scientific objections to the VSM, starting from the records identified by Fathi et al. (2025),[21] recording the queries, the date, the screening rule (counting as scientific only objections of the empirical type, in Schwaninger and Scheef's own distinction) and the exclusions, and publish the resulting bibliography. Price: one systematic scoping review. Any third independent empirical line of objection retires the two-objection charge; finding none strengthens only the count, not the model.

Handover

Nothing in this chapter shows that the VSM is wrong. It shows that we arranged matters so that we would not find out, and that the arrangement was ours to undo at any point. Chapter 29 sets out the two cheapest experiments that would begin to undo it, and Chapter 30 prices the whole repair. The next question is what the neighboring literatures did instead. That is Part Two.

[20]The count is Schwaninger and Scheef's, 6–7 (online-first version): they name the critique "as formulated by Ulrich (1983), e.g.," and the critique "that the model is too general (Amey 1986)," and add that "neither type of criticism is of the empirical type, but rather deductive, and they are often normative, not descriptive." They refer readers to responses in Schwaninger (1989) and Jackson (2000). The charge is made in Ulrich 1981 (see Chapters 8 and 16); Schwaninger and Scheef's count cites it through Ulrich 1983.

[21]Mohammad Reza Fathi, Soraya Birami, Alireza Payvar, and SeyedAli Doorafshan, "A Systematic Review of the Viable System Model: Applications, Insights, and Future Directions," *Journal of Systems Thinking in Practice* 4, no. 3 (2025): 109–45, https://doi.org/10.22067/jstinp.2025.91552.1136. Chapter 14 gives its selection figures.

Part II

Outside: the people who found it without us

The Fight Club

> Judy is older now, and the club is much bigger than her gym. The others came up in different traditions: better funded, better cornered, with sparring records, sponsors and thirty years of published results. Several of them found her best move on their own, gave it another name, and can show what it did. She has never trained alongside any of them. This part introduces her to each in turn.

A bibliography, a factory, a coal face, an assembly line, a paper: most chapters in this part open on one of these. Nothing in the scenes is invented, and each is sourced in a note. The scenes in Part One posed a question by showing a future that had already answered it. These show somebody who reached a piece of this book's argument first, in another vocabulary, and often with better evidence than we have assembled since. They are receipts, and the debt they record is ours.

We come to these literatures as debtors, not competitors. Organization design, hierarchy theory, ecology, sociotechnical systems design and sociocracy, the study of the commons, and holonic manufacturing each converged on components we treat as ours. In most cases they measured what they had, published it, and let others test it; the chapters that follow give the record for each. The part closes with the three published objections that land. They are the ones we have never answered, which are seldom the ones most often quoted at us.

This part asks something different of a reader from Part One, and something harder. Part One asked to be watched. Here the people who use an instrument audit it in use, in that instrument's vocabulary. A reader outside the tradition may skip the terminology and follow the shape of the charge, which is simple: something has been recommended to clients for decades without, as far as we have been able to find, ever being compared with the alternatives on the same evidence. Anybody can assess a charge of that shape. This is the most technical part of the book, and about a third of its length.

There is a reason to do the accounting now. Most of the questions this book raises are already being worked on elsewhere, in other vocabularies, by fields with data, funding and venues, and without the theory. They will be answered either way; what can still go wrong is that they are answered without us, and that the vocabulary, which does hold something the others have not assembled, is retired before anyone establishes what it was worth. Our own literature does not help, for the reason Chapter 9 gives: it has published almost no results against itself.

Chapter 10

Six Components, Six Literatures, No Beer

The second edition of *Brain of the Firm*, published in 1981, ends with a Select Bibliography in four sections. Beer titles the last of them Explicit References. It runs to Ashby, McCulloch, Shannon and Weaver, Forrester, Bremermann, Waddington, Thom, and *The Pauper's Cookbook*, which is there because Beer consulted it often enough to judge that it "must be of some importance," and because it gave the preface to the new edition its last line: "brains need gentle handling or they are apt to disintegrate."

It does not contain Herbert Simon, whose paper on the architecture of complexity is nineteen years old by then. It does not contain Arthur Koestler, who had named the holon fourteen years before. It does not contain Elliott Jaques, whose subject was a near neighbor of the thing the book has just spent its length drawing: levels of work, and how to tell one from another.

Nothing in the list suggests concealment. A man wrote down what he had read. It is the shelf as it stood in the room where the model was assembled, and this chapter is about the rest of it.[1]

Taken by itself, the list is a small piece of evidence about one man's reading. We have made more of the omission than he did. For as long as most of us have taught the model, we have presented its six components (our grouping, not Beer's "six vertical components of corporate cohesion")[2] as though they had a single origin, and we have described that origin to clients and students as a discovery rather than an arrangement. The bibliography shows in print that at least three of the people who found the same components were not on the shelf when the model was assembled. It

[1] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), Select Bibliography, 405–10, with the Explicit References at 410; Beer's reason for the cookery book is given in the Preface to the Second Edition, xii. The cookery book is Jocasta Innes, *The Pauper's Cookbook* (Penguin Books, 1971). The dates are Herbert A. Simon, "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82, and Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967). This is a check of a printed list and is as reliable as such a check can be; it is not a claim that Beer nowhere mentions these men in his running text, although a search of the scanned text of this edition (23 September 2026) found no mention of any of the three, nor of the holon, and the separate list of references for Part Four (397) does not name them either. Chapter 11, The Nesting, uses the same list as evidence and does not repeat the description.

[2] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 217, the First Axiom of Management, which balances horizontal variety against "the six vertical components of corporate cohesion." The six components of this chapter are the authors' grouping of design ideas, not Beer's.

shows nothing about whether they were right, and nothing about whether Beer would have written a different book had he read them. It shows only that the components have more than one address, and that we have been giving out one.

The charge Part One leaves standing

Part One ended with a charge that drawing the model more carefully cannot answer: perhaps the six-part construction is only the vocabulary we inherited. A reader who has come through that hinge is owed an answer, and it cannot come from inside the vocabulary under suspicion.

The objection does not say that the Viable System Model↗[3] is wrong. It says something narrower and more damaging: that the model was never chosen. It was inherited, from a community, from a reading list, from the centenary of a man's birth, and the inheritance has been mistaken for a finding. On this reading the nine chapters of Part One demonstrate nothing about the world. They demonstrate that a sufficiently general vocabulary can be applied to anything, which is a fact about vocabularies and not about organizations.

The charge is serious because it disputes no single chapter. It grants every one of them and asks the one thing none of them answers: what became of the alternatives? If nothing else was ever seriously considered, then the eight defects catalogued in Part One belong to a habit, not to a model that was tested against rivals and kept. A habit does not become a theory by being maintained for a long time.

Beer had an answer of his own, and a concession to go with it. He took the nervous system as his model because it was "a known-to-be-viable system", and held that another, "such as the amoeba", would have served: "The results are the same, as they must be if viability as such has its own laws." In *Decision and Control* he had named two rival sources, the ecological model behind laissez-faire and the neurological model behind the mandatory controller, and rejected both as "manifestly partial."[4] Twenty-three pages after the amoeba he wrote that "the choice of five (rather than some other number) is somewhat arbitrary." He set the standard for the whole construction in a sentence the tradition still quotes, "a model is neither true nor false: it is more or less useful", and tied it to a test the tradition does not quote: if a model's "predictive value is low, we had better start

[3]Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

[4]Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 374–75, paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994): "the doctrinaire laissez-faire man is backing a model of the ecological system, and the doctrinaire controller is adopting a model from the neurological system"; "we want neither of these models; because they are manifestly partial." Beer's alternative is "a model taken from an organic system": "Only a complex creature, such as a man himself, provides an adequate model" (375).

again."[5] None of this meets Part One's charge. An argument from the laws of viability, made by the author of the model, is no independent check on it, and the prediction test Beer named was never run. The outside evidence has to do the work his own argument could not.

Something did become of the alternatives. Several literatures reached pieces of this construction independently, by other doors, and one or two reached them earlier. This chapter owes that accounting: what else is on the shelf, on what evidence each neighbor stands, and what each costs the claim we have been making. Part One has already conceded the deepest of our failures, and nothing here retracts that concession. The inventory is where a defense we can actually support has to start.

The register of this part depends on two things. None of these literatures is an opponent. They are people who found something real without the help of our vocabulary, and the right response is credit, stated at the grain the evidence supports. And the omission on Beer's shelf is far less interesting than the omission on ours. He had a reading list and a deadline. We have had decades, indexes, search engines and a professional literature, and for most of that time, speaking for ourselves, we did not read these neighbors either.

How the neighborhood is sorted

We sort the neighborhood by what each neighbor shares with Beer, not by discipline. Sorted that way, it falls into six groups. Five arrive at a component of the model on their own evidence, and these are less a threat than an argument. The nesting group supplies two components, nesting and criteria for a level, so five sharing groups yield six components. The sixth denies that the job the model sets itself is well-posed at all, and that group is the most likely to appear in a review. The admission rule, and what it costs, is set out in the note near the end of this chapter. Most of the groups below have a chapter of their own after this one; here they get the length an inventory should give them.

Independence, in the entries that follow, means something specific and modest: the earliest statements we could check do not cite Beer or Ashby, and no later acknowledgment of either has come to light. Absence of citation is weaker than proof of ignorance. A researcher can read without citing, and a tradition can carry an idea a long way without naming where it came from. Each finding of independence below is therefore defeasible,

[5] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 75 ("a known-to-be-viable system"; the amoeba) and 98 ("somewhat arbitrary"); Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2 ("So remember: / a model is neither true nor false: / it is more or less useful."); Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113, where the measure of utility "may be accomplished by straightforward testing of the predictive value of the scientific model. If this predictive value is low, we had better start again." *Decision and Control* is paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994), which reproduces the 1966 text.

and the receipt at the end of this chapter prices the work that would firm those findings up or overturn some of them.

Variety matching: Galbraith

Jay Galbraith found that organizational structure is determined by the information that must be processed under uncertainty, so that a designer has two families of moves: reduce the need for information, or increase the capacity to process it. He reached the finding through fieldwork in 1968, set it out in a book in 1973, and compressed it into a paper in 1974. The evidence base is one longitudinal study inside Boeing's commercial-airplane division. That is a narrow base for a general claim, and Galbraith did not present it as wider than it was.

For this chapter the citation record matters. The 1974 paper and the 1968 working paper cite the Carnegie tradition, and neither Ashby nor Beer.

The cost to us is precise. Variety attenuation and amplification were sufficient to describe the problem Galbraith describes, but they were not necessary to see it, because he saw it without them.[6] Chapter 12, The Variety Logic, takes the case in full, including what our search covered and what it does not.

Nesting and levels: Simon, Koestler, Jaques

Three routes reached the nesting. Herbert Simon's near-decomposable hierarchy, in 1962, is a formal argument with a parable attached. Arthur Koestler's holon, coined in 1967 and given its propositions in the same book, is a proposition set. Elliott Jaques's time-span of discretion, developed in a program of organizational measurement from 1956 to 1989, is the only one of the three that tried to make a level measurable rather than asserted. Beer's 1981 Select Bibliography names none of them.

The three evidence bases are unlike one another, and they should be kept apart rather than added together: a formal argument and a parable, a set of propositions, and a measurement program with a mixed empirical record. The cost to us is correspondingly specific. Three independent routes found that wholes nest inside wholes, a strong result for nesting in general and one we did not need to supply. None of them delivered a ruler the others accept, so the nesting is well attested and the criterion for a level

[6] Jay R. Galbraith, *Designing Complex Organizations* (Addison-Wesley, 1973), and "Organization Design: An Information Processing View," *Interfaces* 4, no. 3 (1974): 28–36. The airplane division is *Environmental and Technological Determinants of Organization Design: A Case Study*, Working Paper 352-68 (Sloan School of Management, MIT, 1968). The working paper names the site as "the Commercial Airplane Division of the Boeing Company" in its abstract, as "a longitudinal case study". That he arrived at it without Ashby is a negative finding; Chapter 12 states what was searched and what the finding does not cover. For the 1973 book, a search of the full-text index of the Internet Archive's scan (29 September 2026) finds neither Ashby nor Beer, nor Wiener, Shannon or *cybernetic*; the reference list has not been read page by page.

is not.[7,8,9] Chapter 11, The Nesting: Simon, Koestler, Jaques, follows each route and the state of its evidence separately.

The comparison needs Beer's own criterion, stated once, because every neighbor in this part is measured against it. Beer's recursion is a relation between viable systems, and viable means "able to maintain a separate existence." His theorem reads: "In a recursive organizational structure, any viable system contains, and is contained in, a viable system."[10] He fenced the term on three sides. Recursion is not "any kind of system contained inside another" but "an absolutely precise definition of viability." It is not organizational level: "If the term *recursion* is allowed to become a meaningless appendage to the term *organizational level*, then we do not have a valid model of a viable system." And the managerial parts do not count, since "Systems Two, Three, Four, and Five are not in themselves viable systems at any level of recursion." A layer of management is therefore not a recursion. Beer added that the levels "answer to mathematical criteria", and discussed them "in words and in diagrams" because his was "not a book for mathematicians."[11]

Set against that criterion, the three routes reach objects that are related to Beer's and not the same. Simon's blocks are defined by rates of interaction and need not be viable. Koestler's holon is any part that is also a whole, an organ as much as a firm. Jaques's strata are layers of managerial roles, the reading of level that Beer ruled out. Where their boundaries coincide with his, the accurate description is convergence on a related but not identical object, and whether they coincide is an empirical question that Chapter 11 prices. Koestler's holon later went into factories as holonic manufacturing, whose production record Chapter 15 examines.[12]

---

[7]Herbert A. Simon, "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82. Simon's own retraction of the watchmakers' arithmetic two paragraphs after he gives it, and the near-decomposability definition that does the theoretical work, are quoted in Chapter 11, The Nesting.

[8]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), 48, where the holon is coined, and Appendix I, "General Properties of Open Hierarchical Systems (O.H.S.)," 341–48, where the propositions first appear. The propositions usually quoted come from Koestler's 1969 restatement in the Alpbach appendix, "Some General Properties of Self-Regulating Open Hierarchic Order"; Chapter 11 records which wording we were able to verify and which we were not.

[9]Elliott Jaques, *Measurement of Responsibility: A Study of Work, Payment, and Individual Capacity* (Tavistock Publications, 1956), and *Requisite Organization: The CEO's Guide to Creative Structure and Leadership* (Cason Hall, 1989). The empirical record is mixed and is set out in Chapter 11, including the replications that do not support the criterion.

[10]Beer, *Diagnosing the System*, 1, where "viable" is defined with the *Oxford English Dictionary*; Beer, *Heart of Enterprise*, 118, for the Recursive System Theorem, restated at 308 and in the glossary of rules at 566.

[11]Beer, *Diagnosing the System*, 2; Beer, *Heart of Enterprise*, 308 (organizational level, "mathematical criteria", "since this is not a book for mathematicians, we have to discuss the issues in words and in diagrams") and 411 (Systems Two to Five). Chapters 11 and 14 compare their neighbors with this criterion and do not restate it.

[12]Hendrik Van Brussel, Jo Wyns, Paul Valckenaers, Luc Bongaerts, and Patrick Peeters, "Reference Architecture for Holonic Manufacturing Systems: PROSA," *Computers in Industry* 37, no. 3 (1998): 255–74, for the reference architecture. The count of systems that have run in

Cross-scale coupling: Holling and Gunderson

The panarchy literature describes adaptive cycles coupled across scales through two named and unlike connections, *revolt* and *remember*. Holling's resilience work begins in 1973; the cross-scale account is developed in 2001 and 2002. The dates shape the claim. This group arrives after Beer, so it is no evidence of priority over him. It is evidence of independent rediscovery, a weaker result and still a real one, and the reason the priced receipt at the end of this chapter treats convergence and priority as different findings. The evidence base is an ecological synthesis together with a later review literature.

It costs us a component we draw only in part. Beer's recursion does couple levels in more than one way. Planning, which he called the "glue" of organizational cohesion, runs between recursions; algedonic signals, "alone of all the signals in the enterprise, are organized to break through recursivity"; and System Two and Three-Star "require cross-recursion coordination of a delicate kind."[13] None of these couplings depends on phase, and none of them corresponds to *remember*.[14] Chapter 11 takes the panarchy record as its fourth entry, and Chapter 23 carries the coupling claim forward.

Bounded autonomy: Trist, Herbst, Endenburg

Between 1951 and 1976 the sociotechnical tradition produced semi-autonomous groups and the principle of minimum critical specification. Gerard Endenburg's double-linked circles descend through Kees Boeke's consent principle of 1945. Neither lineage establishes a dependency on Beer, in either direction, and neither claims one. The evidence base is field studies, design principles, and an adoption record that neither community has adequately counted. The missing count is as much our failure as theirs, and we are in a position to remedy it.

This is the most practical entry in the inventory. Beer supplied a rule and mechanisms of his own. The metasystem "should make only that degree of intervention that is required to maintain cohesiveness", and the resource bargain, an audit run "by agreement with System One management" and the sorting of constraints into mandatory and anti-oscillatory are the means.[15]

---

day-to-day production is ours and it is a count of a negative; the three sources it rests on, and what would refute it, are given in Chapter 15.

[13] Beer, *Heart of Enterprise*, 336, for planning as the "glue" of organizational cohesion, and 407, for the algedonic signal; Beer, *Diagnosing the System*, 135.

[14] Lance H. Gunderson and C. S. Holling, eds., *Panarchy: Understanding Transformations in Human and Natural Systems* (Island Press, 2002); C. S. Holling, "Resilience and Stability of Ecological Systems," *Annual Review of Ecology and Systematics* 4 (1973): 1–23, https://doi.org/10.1146/annurev.es.04.110173.000245, where the resilience work begins; C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405. How much of the panarchy literature actually uses the cross-scale couplings, rather than the adaptive cycle alone, is quantified in Chapter 11.

[15] Beer, *Heart of Enterprise*, 158, for the rule, and 180–81, for the red and green paper; Beer, *Diagnosing the System*, 38–39, for the resource bargain, and 82, for the audit.

What we most often supply in practice is a placement: a named box on a diagram. What neither Beer nor we supply is a way to compute the rule, or evidence that the mechanisms bound autonomy as intended. These traditions supply design rules and working mechanisms at that point.[16,17] Chapter 13, The Autonomy Design, sets the rules and mechanisms beside what we offer.

Multi-level viability: Ostrom

Elinor Ostrom's eight design principles describe institutions that persist across generations under conditions theory predicted would destroy them. The coding instrument appeared in 1989 and the synthesis in 1990, and the convergence is independent of Beer though not of Ashby: the Ostroms' multiple levels of rules come, largely through Vincent Ostrom, from Ashby's hierarchy of feedback mechanisms (Chapter 14). Her eighth principle nests enterprises in layers without requiring each layer to be viable, so what she reached is nested governance: related to Beer's recursion and not the same claim. Chapter 14 draws the consequence. The evidence base is a coded case corpus with a published instrument, and the instrument deserves our attention. Because it was published, other researchers could code further cases, disagree in a form that could be checked, and record cases where the principles did not hold.

That is where the cost falls, and it falls on us rather than on the model. This literature measured failures and published disconfirmations. Ours did not.[18] Chapter 14, The Evidence: Ostrom, takes the comparison seriously, including what a comparable instrument for our own claims would have to look like.

---

[16] E. L. Trist and K. W. Bamforth, "Some Social and Psychological Consequences of the Longwall Method of Coal-Getting," *Human Relations* 4, no. 1 (1951): 3–38; F. E. Emery and E. L. Trist, "The Causal Texture of Organizational Environments," *Human Relations* 18, no. 1 (1965): 21–32; P. G. Herbst, *Socio-Technical Design: Strategies in Multidisciplinary Research* (Tavistock Publications, 1974), for "Designing with Minimal Critical Specifications"; Albert Cherns, "The Principles of Sociotechnical Design," *Human Relations* 29, no. 8 (1976): 783–92. Attributing minimum critical specification to the tradition is right at this grain; Chapter 13, The Autonomy Design, narrows it to Herbst and records Cherns's own disclaimer of originality.

[17] Gerard Endenburg, *Sociocracy: The Organization of Decision-Making; "No Objection" as the Principle of Sociocracy* (Eburon, 1998); Kees Boeke, "Sociocracy: Democracy as It Might Be" (May 1945), English text at https://www.sociocracy.info/sociocracy-democracy-kees-boeke/, for the consent principle (original place of publication not identified); A. Georges L. Romme and Gerard Endenburg, "Construction Principles and Design Rules in the Case of Circular Design," *Organization Science* 17, no. 2 (2006): 287–97. Chapter 13 reports that neither the sociocratic movement nor this community publishes an adoption count that would allow a comparison, so both sides of it are unavailable.

[18] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), the eight design principles at 90; Elinor Ostrom, Arun Agrawal, William Blomquist, Edella Schlager, and S. Y. Tang, *Common-Pool Resources (CPR) Coding Manual* (Workshop in Political Theory and Policy Analysis, Indiana University, 1989), for the instrument. The reviews that test the principles against coded cases, and what they found, are in Chapter 14.

The group that denies the job is well-posed

This group will actually appear in a review, and we have engaged it least. Niklas Luhmann's social systems are operationally closed and constituted by communication, which makes an organization something that cannot be steered from outside at all.[19] Assemblage theory, in Manuel DeLanda's rendering, insists that a component can be detached from one whole and plugged into another without ceasing to be itself.[20] Actor-network theory abolishes levels entirely.[21] Werner Ulrich's critical heuristics asks, of any systems study, whose boundary judgments are in force and who is affected but not consulted.[22] Michael Jackson's classification places the VSM among the functionalist methodologies, useful where interests are unitary and untrustworthy where they are coercive.[23] Hayek's argument that no regulator can assemble the knowledge dispersed among the regulated is older than most of it, and Chapter 16 finds one layer of it, the measure problem, still unanswered.[24] And process philosophy, from Whitehead to Robert Cooper, holds that an entity is an artifact of a description that moves more slowly than the thing described.[25] Chapter 16, The Objections That Land, states each objection at the strength its author would want and asks which of them land.

---

[19] Niklas Luhmann, *Social Systems*, trans. John Bednarz Jr. with Dirk Baecker (Stanford University Press, 1995), and *Organization and Decision*, ed. Dirk Baecker, trans. Rhodes Barrett (Cambridge University Press, 2018). Chapter 16, The Objections That Land, states the objection at the strength Luhmann would want and says what survives it.

[20] Manuel DeLanda, *A New Philosophy of Society: Assemblage Theory and Social Complexity* (Continuum, 2006), 10–11, for relations of exteriority. Chapter 16 takes the argument up, together with Beer's own statements on detachment.

[21] Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford University Press, 2005), Part II, "How to Keep the Social Flat," from 165, with "keep the social domain completely flat" at 171 and "the Flat-Earthers of social theory" at 172. The application to Cyberstride and the operations room — which is ours and not Latour's — is in Chapter 16.

[22] Werner Ulrich, *Critical Heuristics of Social Planning: A New Approach to Practical Philosophy* (Paul Haupt, 1983), and "Critical Heuristics of Social Systems Design," *European Journal of Operational Research* 31, no. 3 (1987): 276–83. His 1981 critique of the Chilean project, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59, is read from the source in Chapter 16, and it is the hardest thing in this book to answer.

[23] Michael C. Jackson and Paul Keys, "Towards a System of Systems Methodologies," *Journal of the Operational Research Society* 35, no. 6 (1984): 473–86; Michael C. Jackson, *Systems Thinking: Creative Holism for Managers* (John Wiley & Sons, 2003). What the 1984 grid does and does not classify — coercive contexts are named there and deliberately left outside it — is set out in Chapter 16.

[24] F. A. Hayek, "The Use of Knowledge in Society," *American Economic Review* 35, no. 4 (1945): 519–30. Chapter 16 separates the objection into layers and says which of them the model can meet and which it cannot.

[25] Alfred North Whitehead, *Process and Reality: An Essay in Cosmology* (Macmillan, 1929), for the position, and *Science and the Modern World* (Macmillan, 1925), where the fallacy of misplaced concreteness is named. The management-studies form, Robert Cooper's among it, is cited in Chapter 16, whose note on Tsoukas and Chia records that they are not Whiteheadians.

What the convergence establishes

The five sharing groups answer the charge of arbitrariness, and the form of that answer is easy to overstate. It is not that the VSM is confirmed by the company it keeps; resembling its neighbors makes no construction true. It is that every one of its components has been reached independently by someone who had not read Beer: variety matching by Galbraith, nesting by Simon and Koestler, criteria for managerial strata by Jaques, cross-scale coupling by Holling, bounded autonomy by Trist, nested governance by Ostrom. Four of those entries are convergence on a related but not identical object. The nesting of Simon and Koestler, Jaques's strata and Ostrom's nested enterprises do not require each nested layer to be a viable system, as Beer's recursion does, and Holling's coupling carries a phase that Beer's couplings lack. A construction assembled from parts that six other literatures found, or found versions of, on their own may still be wrong. It is not arbitrary.

Inside that answer sits an asymmetry we have not always observed. Where a neighbor arrived earlier, the accurate word is priority, and priority belongs to them. Where a neighbor arrived later by an unrelated route, the accurate word is convergence, which strengthens the case for the component and settles nothing about who found it first. We have at times used the second sort of evidence as though it were the first, and the panarchy entry is where the temptation is strongest.

The answer has a price, and Part Two hands it to Part Four. If each component has an independent address, we cannot say that Beer discovered the components, and we have said it, in teaching and in print, for a long time. What survives is narrower: these components were assembled into one construction, with couplings specified between them, and nothing else on this shelf provides that assembly. It is an assembly claim, not a discovery claim. Chapter 23, The Narrow Claim, states it precisely and tests whether it holds. This chapter makes the credit real and hands the narrowed claim forward without inflating it on the way.

How we looked

We assembled the neighborhood above by asking what each candidate shares with Beer rather than which discipline it belongs to, which is why the groups cut across management studies, ecology, sociology, philosophy and economics. A candidate was admitted if it makes a claim about the same object, the persistence of a bounded system under complexity it cannot absorb, and left out if it only shares vocabulary. That rule is ours, and it is contestable: it admits Friston[26] and Ostrom and excludes most of the strategy literature.

[26] Karl Friston, "The Free-Energy Principle: A Unified Brain Theory?," *Nature Reviews Neuroscience* 11, no. 2 (2010): 127–38, https://doi.org/10.1038/nrn2787; Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo, "Active Inference: A Process Theory," *Neural Computation* 29, no. 1 (2017): 1–49; Thomas Parr, Giovanni Pezzulo, and Karl J. Friston, *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (MIT Press, 2022). The technical criticism of the principle, which is substantial, is in Chapter 12.

Two disclosures. We are inside the community this chapter is about, and we are not neutral about the outcome. And this is a shelf we can reach. The works here are in English or in translation into it, and a literature that never crossed over would be invisible from where we stand, which is exactly how a literature that does not exist would look. The Select Bibliography check was repeated on 23 September 2026 against a copy of the 1981 edition; the cross-literature checks were last repeated on 12 August 2026.

The priced receipt

Receipt — completeness and priority audit of the six components. Commission an audit that identifies, for each component, the earliest independent statement, and checks the citations in both directions. One researcher, six bounded searches, a preregistered evidence table, and external spot-checking: approximately fifteen working days. The evidence table records for each entry whether the object reached is Beer's (a viable system within a viable system) or a related one; Simon and Koestler, Jaques, Ostrom and Holling are entered as related, so a negative result on them narrows the claim about nesting, not about recursion. A negative result — that one or more components cannot be shown to have an independent lineage — would narrow the non-arbitrariness claim to the components that survive, and we would print the narrowed version. Holling is already the warning: convergence after Beer is evidence of independent rediscovery, not priority.

Chapter 11

The Nesting: Simon, Koestler, Jaques

Chapter 10 opened on the Select Bibliography of *Brain of the Firm*; it has no room for Simon, Koestler or Jaques.[1] The earlier *Decision and Control* names none of the three either.[2] The name Simon does occur once in *Brain of the Firm*. It belongs to Beer's son, Simon Beer, the electronic engineer who built in England the prototype set of algedonic meters that was taken to Chile.[3]

The secondary literature is just as quiet. Andrew Pickering's *The Cybernetic Brain*, a scholarly history of British cybernetics, has no Koestler in its index and one incidental Simon, on a page about the Logic Theorist.[4] The silence runs the other way too. Simon's "The Architecture of Complexity" cites Wiener and Ashby and names neither Beer nor Koestler, and the third edition of *The Sciences of the Artificial* mentions none of them.[5] One link runs between any two of these bodies of work, and only one: Koestler read

[1]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), Select Bibliography, 405–10. The named works are Beer's own listing; *The Pauper's Cookbook* is Jocasta Innes, *The Pauper's Cookbook* (Penguin Books, 1971), entered under "D. Explicit References" (p. 410) and quoted in the preface to the second edition (p. xii): "brains need gentle handling or they are apt to disintegrate". The absence of Simon, Koestler and Jaques from the Select Bibliography is a check of a printed list; a full-text search of the same edition, run on 23 September 2026, finds none of the three in the running text either.

[2]Stafford Beer, *Decision and Control: The Meaning of Operational Research and Management Cybernetics* (John Wiley & Sons, 1966). The negative claim is ours. It rests on a full-text search, run on 23 September 2026, of the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994), which reproduces the 1966 text: neither Simon, Koestler nor Jaques is named anywhere in the book, index and apparatus included. A search of a digital text layer is as good as the text layer, and we report it as that.

[3]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 284: "A prototype system of ten algedonic meters, linked by a single wire in a loop through a large summation meter, was built in England by my son Simon Beer (an electronic engineer), and taken to Chile for experimental purposes." The book's index has the single entry "Beer, Simon, 284", and a full-text search of the edition finds no other occurrence of the name.

[4]Andrew Pickering, *The Cybernetic Brain: Sketches of Another Future* (University of Chicago Press, 2010). The claim rests on the book's index. We were unable to run a full-text search of the volume, and the reader should treat "no Koestler at all and one incidental Simon" as an index result rather than a concordance.

[5]Herbert A. Simon, "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82. Verified in full text: the paper's notes cite Shannon and Weaver and W. R. Ashby's *Design for a Brain* and "Design for an intelligence amplifier," and Wiener once; Beer, Koestler and Jaques do not appear (search of the text layer of the JSTOR scan, repeated on 29 September 2026). For the third edition, Herbert A. Simon, *The Sciences of the Artificial*, 3rd ed. (MIT Press, 1996); we have not been able to search that volume in full text and rest the claim on its index and bibliography.

Simon. He says so on the first page of the chapter in which he coins the holon, where the parable of the two watchmakers comes with the words "I owe it to Professor H. A. Simon", and he credits him again in the book's appendix, for the word *span*.[6]

The five do not form a tradition. Four people besides Beer arrived at the nesting of systems within systems, and they arrived separately. Three of them were working within about fifteen years of one another and of Beer: Jaques in 1956, Simon in 1962, Koestler in 1967. Two of them worked in the same country and in overlapping industrial worlds. The fourth, C. S. Holling, came after Beer's first statement of the model, so his work is convergence rather than prior art, and for that reason a different kind of evidence. Independent arrival is stronger news than influence would have been. When five people working from ecology, from economics and administration, from psychoanalytic industrial fieldwork, from evolutionary biology and from operational research all end up describing systems nested inside systems, the structure is unlikely to be an artifact of any one of their vocabularies.

That is good news for the model. The less comfortable news is why the four entries below share one form. Each arrival came with an account of what a level is. None came with a way of measuring where one level ends and the next begins. Simon has a threshold he cannot set. Koestler has propositions he cannot test. Jaques has a genuine ruler whose construct validity his own side concedes is thin. Holling has a phase nobody has shown they can identify in the present tense. And Beer, as Chapter 5 argued at length, has a criterion for a level of recursion that resolves in practice to a judgment call.

Chapter 5, *The Ladder Without a Ruler*, made that charge against our own side and named these neighbors in passing. Here the neighbors get their own hearing: what each of them claimed, on what evidence, with what admitted limits, and what the VSM↗[7] would have gained by reading them. Where the material overlaps with Chapter 5 we send the reader there and go further here. Panarchy has a place in a chapter named for three earlier authors because the question in this part of the book goes beyond who

[6]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), 45 for the parable ("I owe it to Professor H. A. Simon"), 47 for Simon's conclusion on stable intermediate forms, quoted at length, 52 n. for Simon's decomposable hierarchies, 342 for proposition 2.2 with *span* credited to Simon, and 364 for "The Architecture of Complexity" in the bibliography (read in the Internet Archive scan of the Hutchinson first edition on 2 October 2026; page 45 checked against the page image). The Alpbach appendix repeats the credit: Arthur Koestler, "Beyond Atomism and Holism—The Concept of the Holon," in *Beyond Reductionism: New Perspectives in the Life Sciences*, ed. Arthur Koestler and J. R. Smythies (Hutchinson, 1969), appendix, "Some General Properties of Self-Regulating Open Hierarchic Order," proposition 2.2: "The number of levels which a hierarchy comprises is a measure of its 'depth', and the number of holons on any given level is called its 'span' (Herbert Simon)." The link is a debt declared in the book that coins the holon, not only an attribution in the technical apparatus. We found no reciprocal citation, and no citation in either direction between any other pair among the five.

[7]Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

named hierarchy first. It asks who made cross-scale nesting operational, and how far they got.

One caution applies to all four entries. Beer's criterion, set out in Chapter 10, is a criterion for a recursion and not for a level in general: a viable system inside a viable system, which is neither an organizational level nor a layer of management.[8] Simon's blocks, Koestler's holons, Jaques's strata and Holling's scales are each a related object. Where their boundaries meet Beer's, the finding is convergence on a related but not identical object, and whether they meet at all is what the receipt at the end of this chapter tests.

How we looked

Several of this chapter's claims are claims about an absence, and an absence is the easiest thing in scholarship to get wrong. Where a full-text search was possible we ran one: Simon's 1962 paper and his 2002 one, and Koestler's Alpbach appendix. Where only an index was available we say so. *Decision and Control*, the second edition of *Brain of the Firm* and *The Heart of Enterprise* we searched in full text; the Pickering result is an index result rather than a concordance. The Select Bibliography we checked as a printed list and then against the running text of the book, where none of the three men is named either.

The same caution applies to the count of holonic systems that reached day-to-day production, which this book makes in Chapter 15 rather than here. That count rests on three sources that would have reported a deployment if one existed, and it is exactly the kind of claim we are asking other people to make about us. One counterexample retires it, and we would like to be sent it.

Simon — a level is a rate separation

What he found · Near-decomposability: interactions are stronger or faster within a level than between levels. **When** · 1962, revisited in 2000–2002. **Who cited whom** · The searches reported above found no Beer–Simon exchange on this claim. **Evidence base** · An analytical argument whose watchmaker arithmetic is withdrawn two paragraphs later. **What it costs us** · The epsilon that distinguishes a level remains unset.

"The Architecture of Complexity" was read to the American Philosophical Society and published in December 1962, ten years before *Brain of the Firm*.[9] Most readers remember it as the parable of the two watchmakers, and remember the parable in a way that misses what Simon was doing with it.

Hora builds in subassemblies of ten, three levels deep. Tempus builds each watch of a thousand parts in one sitting and loses everything whenever

[8] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 1–2; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 118, 308 and 411. The quotations are given in Chapter 10.

[9] Simon, "The Architecture of Complexity," 467–82. Read to the American Philosophical Society and published in the issue dated 12 December 1962.

the telephone rings. At an interruption probability of one in a hundred, Tempus takes about four thousand times as long. Two paragraphs later comes the comment almost nobody quotes: "These assumptions are not particularly realistic... Although we cannot, therefore, take the numerical estimate seriously the lesson for biological evolution is quite clear and direct. The time required for the evolution of a complex form from simple elements depends critically on the numbers and distribution of potential intermediate stable forms."[10] The four thousand is a rhetorical prop, and Simon says so in the same breath. The theorem is the final sentence. It is a claim about the availability of stable intermediates, and ratios play no part in it.

The definition underneath the parable matters more for us. Simon defines hierarchy by *near-decomposability*, not by authority, and he states the exclusion outright: "We shall want to include systems in which there is no relation of subordination among subsystems. (In fact, even in human organizations, the formal hierarchy exists only on paper...)"[11] A level exists where interaction rates separate: the parts within a block interact with each other much faster, or much more strongly, than the blocks interact with one another. Any reader of Beer will recognize the epistemological payoff: "we can study the equilibrium at any given level without much concern for the slow dynamics at the levels above, and we can study the dynamics at any level without much concern for the rapid return to equilibrium at the levels below."[12] That is the license to model one level at a time, and the VSM exercises it on every page.

Beer spent his career insisting that recursion was not the organization chart. Simon had got there in 1962, from a different direction, and with something Beer claimed and never supplied: a criterion that is in principle computable. Beer wrote that the levels of recursion "answer to mathematical criteria" and then, since his was "not a book for mathematicians",

[10] Simon, "The Architecture of Complexity," 470 for the watchmakers and the four-thousand estimate, 471 for the retraction: "These assumptions are not particularly realistic. They undoubtedly underestimate the decrease in probability of achieving the assembly with increase in the size of the assembly... Although we cannot, therefore, take the numerical estimate seriously the lesson for biological evolution is quite clear and direct." The ellipsis in the text stands for a sentence in which Simon says the assumptions understate the advantage of hierarchy.

[11] Simon, "The Architecture of Complexity," 468. Simon's parenthesis reads in full: "(In fact, even in human organizations, the formal hierarchy exists only on paper; the real flesh-and-blood organization has many inter-part relations other than the lines of formal authority.)"

[12] Herbert A. Simon, "Near Decomposability and the Speed of Evolution," *Industrial and Corporate Change* 11, no. 3 (2002): 587–99. The sentence quoted is from the later paper, not from "The Architecture of Complexity" — we searched the 1962 text and it is not there. It reads in full: "In fact — and this property is used continually by the sciences — we can study the equilibrium at any given level without much concern for the slow dynamics at the levels above, and we can study the dynamics at any level without much concern for the rapid return to equilibrium at the levels below, or for the (very slowly moving) disequilibrium states at the levels above." We verified the wording against Simon's own working paper, Complex Information Processing Working Paper #549, dated 4 April 2000, in the Carnegie Mellon collection of his papers.

discussed them in words and diagrams.[13] An interaction matrix either has block structure or it does not, and whether it does depends on observable rates, not on the analyst's reading of a system's identity. Had the two literatures met, the obvious joint project would have been to test whether the boundaries a VSM diagnostician draws coincide with the blocks a rate matrix finds. As far as we can establish, that project was never proposed, let alone run.

In principle carries a great deal of weight there, and Simon knew it. His own admissions are more damaging than most of his critics'. Near-decomposability is "a rather strong property for a matrix to possess, and the matrices that have this property will describe very special dynamic systems — vanishingly few systems out of all those that are thinkable." He defends its usefulness on purely empirical grounds: the world we happen to live in is full of such systems. Worse for our purposes, the property is threshold-relative, and the analyst chooses the threshold: "How few they will be depends, of course, on how good an approximation we insist upon." Simon offers no principled way to set it. When he applies the idea to a whole economy, he has to qualify it at once. The input-output matrix is nearly decomposable "with one qualification. There is a consumption subsystem of the economy that is linked strongly to variables in most of the other subsystems."[14] One hub, and the property degrades. The networks organizations actually run on are made of hubs: shared services, enterprise systems, a finance function that touches everything.

The criterion is therefore computable and underdetermined at once, which is a more interesting position than being simply right or simply wrong. It turns a question about identity into a question about a number, and then leaves the number to judgment. Beer's defect is a different one, and larger. Two analysts who disagree about an epsilon can show each other their matrices and see exactly where they part company. Two analysts who disagree about whether a unit has an identity of its own have little to show.

There is a second Simon paper, and keeping the two apart is the most useful thing anyone can do with him. In the work published in 2002 he conceded that the 1962 composition argument does not carry over to multicellular life. He takes bacterial complexity as its limit and writes that existing multi-celled organisms are unlikely to have been formed by the combination of single-celled ones. In its place he built a different argument: among systems of comparable complexity, the nearly decomposable ones adapt faster, because feedback on a part is not scrambled by everything else the system is doing.[15] The claim is better, closer to what a management

[13] Beer, *Heart of Enterprise*, 308.

[14] Simon, "The Architecture of Complexity," 475, for all three quotations: the "vanishingly few" concession, "How few they will be depends, of course, on how good an approximation we insist upon," and the consumption-subsystem qualification.

[15] Simon, "Near Decomposability and the Speed of Evolution," 587–99. The concession is explicit: having taken bacterial complexity as the limit of the composition argument, Simon

cybernetician would want, and much less famous. It concerns the rate of adaptation rather than the possibility of assembly, which makes it testable in settings where assembly arguments are not.

What would settle Simon's side is a demonstration, on real organizational interaction data, that the block structure is robust across a reasonable range of thresholds, or that it is not. The data required (who talks to whom, how often, about what) is the kind most organizations now generate as a byproduct of their own communications infrastructure. That is a change in the world since 1962, and this field has not exploited it.

Koestler — the word, and the propositions

What he found · The holon and a proposition set for self-regulating hierarchy. **When** · 1967, restated in 1969. **Who cited whom** · Koestler cites Simon, for the watchmakers' parable and for "span"; Beer and Koestler do not cite one another in the material checked. **Evidence base** · About sixty-five propositions, not a validated measurement program. **What it costs us** · A better vocabulary and explicit channel rules still leave the claims untested.

The Ghost in the Machine appeared in 1967, and Koestler's first contribution was a name: the holon, the unit that is a whole to what it contains and a part to what contains it. "Parts and wholes in an absolute sense do not exist in the domains of life."[16] Beer's noun for the thing that recurses is "viable system", and he used the mathematical term itself as a count noun: the embedded systems "are themselves viable systems. They are RECURSIONS OF THE VIABLE SYSTEM."[17] That noun names the recursing thing only if it is viable. Koestler's holon names any part that is also a whole, which is wider than Beer's term and easier to use, and the difference matters more than style. It is one reason VSM diagrams are hard to read aloud and holonic ones are not. It is also one reason our diagrams are so often mistaken for organization charts, since the nouns practitioners reach for to label them are the nouns of the org chart. Naming was Koestler's contribution, and a considerable one.

The technical apparatus came in the same book, as its first appendix: propositions on the "General Properties of Open Hierarchical Systems", across ten headed sections. Two years later Koestler restated them, some

writes that "it is highly unlikely that existing living multi-celled organisms were formed by the combination of single-celled organisms," and builds the fitness-rate argument in its place. On his age: the paper appeared posthumously in 2002, Simon having died in February 2001; the working-paper version is dated 4 April 2000, when he was eighty-three. He turned eighty-four in June 2000. We have not been able to date the final revision and therefore cannot confirm "at eighty-four" as more than approximate.

[16] Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), 48, where the holon is coined. The sentence quoted is proposition 1.3 in both of Koestler's proposition sets, and we give it in the form of the 1969 Alpbach appendix. The 1967 book prints it in its own Appendix I, 341, with "domain" for "domains" (read in the Internet Archive scan of the Hutchinson first edition on 2 October 2026; proposition 1.3 checked against the page image).

[17] Beer, *Diagnosing the System*, 2.

sixty-five in all, in the appendix to his Alpbach paper, and the quotations below follow that version.[18] Anyone trained in this model who reads that appendix beside *The Heart of Enterprise* has a strange experience.

Koestler has the autonomy-cohesion↗ tension, and states it as a dual tendency belonging to every holon rather than as a management problem: "Every holon has the dual tendency to preserve and assert its individuality as a quasi-autonomous whole; and to function as an integrated part of an... larger whole." He has attenuation↗ and amplification without either word: "Output hierarchies spell, concretize, particularize. Input hierarchies digest, abstract, generalize", with filters on every echelon stripping noise from the input on its way up. He has the prohibition on skipping levels, stated as a structural rule rather than as etiquette: "The higher echelons in a hierarchy are not normally in direct communication with lowly ones, and vice versa; signals are transmitted through 'regulation channels', one step at a time". He also names the pathology that follows from short-circuiting them. He has recursion as representation, the formally exact version of the idea: a holon is "a system of relata which is represented on the next higher level as a relatum." And he has the answer to the objection every VSM presentation still meets, that hierarchy must mean a tree. Hierarchies, in his account, are " 'vertically' arborizing structures whose branches interlock with those of other hierarchies at a multiplicity of levels and form 'horizontal' networks."

All of that was in print in 1967, five years before Beer set out the same structure in *Brain of the Firm* and twelve before *The Heart of Enterprise*, and neither man appears to have known of the other's work.[19] Several of the pieces were Beer's before they were Koestler's: *Decision and Control*, in

[18]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), Appendix I, "General Properties of Open Hierarchical Systems (O.H.S.)," 341–48, in ten headed sections from "The Janus Effect" to "Regeneration"; the propositions quoted in this paragraph appear there as 4.1 (343), 5.6 and 5.8 (344), 5.11 and 6.1 (345) and section 7 (346), with the short-circuiting pathology as 7.2, in British spelling ("concretise, particularise", "arborising") (read in the Internet Archive scan of the Hutchinson first edition on 2 October 2026). We quote the 1969 restatement: Koestler, "Beyond Atomism and Holism—The Concept of the Holon," appendix, "Some General Properties of Self-Regulating Open Hierarchic Order." The propositions run to about sixty-five across ten headed sections; the ones quoted in this paragraph are, in order, 4.1 (the dual tendency), 5.11 (output hierarchies spell, input hierarchies digest) with 5.8 for the filters that strip the input of noise, 7.1 (regulation channels, one step at a time) with 7.3 for the pathology of short-circuiting them, 5.6 (a system of relata represented on the next higher level as a relatum), and 6.1 (arborization and reticulation). The paper was given at the Alpbach Symposium in 1968 and published in 1969; a version of the essay was also reprinted as Arthur Koestler, "Beyond Atomism and Holism—the Concept of the Holon," *Perspectives in Biology and Medicine* 13, no. 2 (1970): 131–54. We verified the propositions against a full transcription of the appendix and cite them by number because we could not confirm page numbers in the Hutchinson printing.

[19]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), and *Brain of the Firm* (Allen Lane, 1972), the first edition; the second edition of 1981 is the one whose Select Bibliography is described above. The claim that neither man knew of the other is a claim about citation, not about acquaintance: neither cites the other, and we found no correspondence or third-party report that they had read one another. A full-text search of *The Heart of Enterprise* and of the second edition of *Brain of the Firm* finds no mention of Koestler or of the holon.

1966, already has subsystems of greater and lesser autonomy, a many-one reduction of variety through an echelon of small filters, and an algedonic loop that can short-circuit the whole machinery.[20] What Koestler had first is the proposition set and the name. That is a claim about citation, and it goes no further: neither cites the other, and we found no correspondence or third-party report that they had read one another. One letter in one archive would overturn it.

What Koestler lacks is a theory, and the gap can be stated exactly. The propositions predict nothing and forbid nothing. Between them, the self-assertive and the integrative tendencies can be made to account for any observed behavior of any part of any system, which amounts to accounting for none. Nothing in the appendix would tell an observer that a putative holon is not one, and no quantity in it could come out wrong. The VSM has the same difficulty, here reached by a writer with no organizational data at all, and the resemblance runs all the way down: a rich generative vocabulary, immediate intuitive assent from readers, and nothing that could fail.

The container has aged badly as well. *The Ghost in the Machine* argues that the human brain is an evolutionary botch and that the species may need chemical correction. The holon survived its vehicle, which is more than the vehicle deserved, and the survival is instructive. What carried was the part that named a structure people could see in their own material. The claims about the world were left behind.

What would settle Koestler's side is narrower than it looks. Two of his propositions are close to testable as they stand. The rule that signals pass one step at a time, together with the pathology he attributes to short-circuiting it, is a claim about observable communication traffic and observable consequences. Anyone with a year of routing data and a record of escalations could put it at risk. Nobody in our tradition has done so, in the half-century during which we have taught the same rule as the algedonic↗ exception, and that says something about us rather than about him.

Jaques — a ruler

What he found · Time-span of discretion as an empirical criterion for organizational strata. **When** · 1956–1989. **Who cited whom** · Beer and Jaques worked in adjacent British industrial worlds without a citation trail between them. **Evidence base** · Mixed replications and an advocate bibliography that records its own gaps. **What it costs us** · The available ruler has weak construct-validity evidence.

[20] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 104–5 (control systems in the enterprise and the body are both hierarchical and redundant, and "both incorporate subsystems of greater and lesser autonomy"), 384 (the "many-one transformation" by "an echelon of small, rather simple black boxes"), 388 (the importance filter) and 396 (the algedonic loop, "which can effectively short-circuit the total machinery"), paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994), which reproduces the 1966 text.

Elliott Jaques worked the same ground as Beer, in the same country, in the same decades: British industry, the structure of the firm, the shape of authority inside a manufacturing company. His fieldwork was at the Glacier Metal Company while Beer was head of operational research at United Steel. *Measurement of Responsibility* appeared in 1956, three years before *Cybernetics and Management*.[21] Beer never cites him. Of the documented silences in this chapter, this is the one we find hardest to account for, because it is the only one that distance of discipline or of geography cannot explain. Two men were asking what distinguishes one level of a firm from the next, in different senses of the word level, in the same industrial economy, at the same time, and the record we can reach shows no contact.

Jaques claims that the level of work in a role is measurable by a single quantity, the time-span of discretion: the longest period over which a subordinate exercises discretion before the results of that discretion become reviewable. He also claims that organizational strata are discrete rather than continuous, with boundaries falling at roughly three months, one year, two years, five, ten and twenty.[22] Set beside Beer's criterion, the difference is not one of rigor or ambition but of what the criterion produces. Beer's produces a judgment about a system's identity. Jaques's produces a number attached to a role, obtained by asking a manager a question about a task. Two analysts who disagree about whether a unit is viable in its own right and one of the units that do what the containing system does, which is Beer's test in our paraphrase, have no procedure for settling it, and Beer supplies none.[23] Two analysts who disagree about a time-span can go and ask.

Beer did leave a rule of thumb for a missing level: "up to seven or eight: not many more, I hope, or else you may be missing a whole level of

[21] Elliott Jaques, *Measurement of Responsibility: A Study of Work, Payment, and Individual Capacity* (Tavistock Publications, 1956); Stafford Beer, *Cybernetics and Management* (English Universities Press, 1959). Jaques's fieldwork was at the Glacier Metal Company, reported in Elliott Jaques, *The Changing Culture of a Factory* (Tavistock Publications, 1951); Beer was head of operational research at United Steel over the same years: "In the mid 1950s, I was directing management sciences for the largest steel company in Britain, United Steel", in a department called "The Department of Operational Research and Cybernetics" (Stafford Beer, *Brain of the Firm*, 2nd ed. [John Wiley & Sons, 1981], 384). That Beer does not cite Jaques rests on full-text searches of *Decision and Control*, *Brain of the Firm* (2nd ed.) and *The Heart of Enterprise*; that Jaques does not cite Beer we take from Jaques's reference apparatus.

[22] Jaques, *Measurement of Responsibility*, for the derivation of time-span of discretion, and Elliott Jaques, *Requisite Organization: The CEO's Guide to Creative Structure and Leadership* (Cason Hall, 1989), for the stratum boundaries at three months, one year, two years, five, ten and twenty years in their settled form. The lowest five of those boundaries are the ones Rowbottom and Billis test below; the twenty-year boundary is not independently corroborated in anything we found.

[23] The wording is our paraphrase of Beer and not a quotation. His nearest statements: the operational elements of a viable system exhaust its "basic activities, namely those which exist to do what the system does" (Beer, *Heart of Enterprise*, 116); a division that "could in principle be hived-off as a separate business" is a viable system, and the accounting system that prepares the group accounts is not (ibid., 118); and "the viable system is necessarily a producer of the organization, and not just an adjunct to it, however important" (Beer, *Diagnosing the System*, 8).

recursion." In *The Heart of Enterprise* he also published an identity exercise, marked on cards, of which he wrote that "maybe (one day) thresholds will be established, to exceed which will predict collapse."[24] Neither locates a boundary. The second is the nearest thing to a ruler in Beer's own work, and nobody in the tradition built on it.

A ruler does not make a theory right, and the evidence for Jaques's ruler sorts into three unequal piles.

The strongest pile supports the claim Jaques started from, which concerns felt-fair pay rather than structure. Roy Richardson's Minnesota doctorate found a correlation of about 0.86 between the time-span of discretion of a role and what incumbents independently judged to be fair pay for it, on 180 managers at three Honeywell facilities.[25] That is a serious number, obtained from outside the movement, on a sample large enough to take seriously. The movement's own bibliography records its limitation: neither the dissertation nor the book contains the raw data or a graph, so the figure has to be taken as reported.

The middle pile is the discreteness of the strata. This is the part the VSM would actually want, since a boundary is what we are short of. Here the support comes largely from Jaques's own circle and is reprinted in his own edited volumes, which makes it a research program citing itself rather than an independent replication. The one genuinely independent corroboration we found is Rowbottom and Billis. In 1977, from action research in social services departments in England and Wales, they derived five qualitative work strata and reported that the boundaries coincided with Jaques's.[26] The word independent needs one qualification, and we would rather enter it ourselves. Rowbottom and Billis worked inside Jaques's framework of strata while dropping time-span of discretion as the measure.

[24] Beer, *Diagnosing the System*, 19; Beer, *Heart of Enterprise*, 417–33, with the sentence quoted at 432. Beer introduces the scoring with a caution of his own: "we are entering ground where it would be bogus to pretend that proper measurements can be made. These are subjective judgments for which we ask" (419).

[25] Roy Richardson, "An Empirical Study of Fair Pay Perceptions and Time Span of Discretion" (PhD diss., University of Minnesota, 1969), published as *Fair Pay and Work: An Empirical Study of Fair Pay Perception and Time Span of Discretion* (Southern Illinois University Press, 1971). The correlation is +0.862, significant at the 0.001 level, on 180 managers at three Honeywell facilities. We take the figure from Ken Craddock, *Requisite Organization Annotated Bibliography on Elliott Jaques, Wilfred Brown, and Requisite Leadership (Covering 1942 to 2009)*, 5th ed. (August 2009), https://globalro.org/system/files/documents/RO-Biblio.pdf, which reports it from the dissertation and the book; neither contains the raw data or a graph, a limitation the bibliography itself notes.

[26] Ralph Rowbottom and David Billis, "The Stratification of Work and Organizational Design," *Human Relations* 30, no. 1 (1977): 53–76. Five successive levels derived from action research in social services departments in England and Wales — prescribed output work, situational response work, systematic service provision work, comprehensive service provision work, comprehensive field coverage work — with the paper stating that Jaques's "critical time-span boundaries between strata — 3 months, 1 year, 2 years, 5 years, and 10 years — will correspond to the boundaries between successive strata" they identify. One qualification on the word *independent*: Rowbottom and Billis used Jaques's framework of strata while dropping time-span of discretion as the measure, so this is convergence from a different method inside a related research tradition, not corroboration by strangers.

That is convergence from a different method within a related tradition, short of corroboration by strangers. It is still the best thing in this pile.

The thinnest pile is the later work on cognitive maturation curves. We could find no independent psychometric validation of it at all, and the movement's own annotated bibliography points to no such study.[27] That is a negative result of our own searching, and the absence of a study we could find does not prove that none exists.

Three lines of criticism stand against the theory. Only one of them is fatal to the use we would want to make of it.

Alan Fox's 1966 appraisal makes the deepest objection, and it concerns pluralism rather than measurement. On Fox's account a firm is a coalition of divergent interests rather than a single team, so a norm of equitable payment is the residue of a power relation, not a fact waiting to be discovered. On that reading, a correlation between time-span and felt-fair pay may show only that both quantities track an internalized and contested status order.[28] We were not able to obtain the thirty-page pamphlet itself and read Fox's argument through the accounts of others. The characterization above is therefore a reconstruction from his position elsewhere in the same year, not a report of the appraisal's own words.

Writing in *Personnel Psychology* in 1969, Michael Gordon charged the theory with unfalsifiability: it leans on unconscious concepts that can be used to nullify what people actually say about their own work. He proposed three concrete validation routes. Decades later, the movement's own annotated bibliography concedes that they "have yet to be explored".[29] Paul Goodman's 1967 study reported low correlations between time-span and hierarchical level. The movement's bibliography disputes the result on two grounds, and we report it without adjudicating.[30]

---

[27]The negative finding is ours and it is a negative: searching the psychometric literature we found no independent validation of the cognitive-maturation curves, and the annotated bibliography of the movement itself points to no such study. Absence of a result we could find is not proof that none exists.

[28]Alan Fox, *The Time-Span of Discretion Theory: An Appraisal*, Research Paper No. 1 (Institute of Personnel Management, 1966), 30 pp. Fox's pluralist position — that a firm is a coalition of divergent interests rather than a single team, so that norms of equitable payment express a power relation — is set out in the same year in his *Industrial Sociology and Industrial Relations*, Research Paper 3 for the Royal Commission on Trade Unions and Employers' Associations (HMSO, 1966). We were not able to obtain the thirty-page appraisal itself and read Fox's argument in it through the accounts of others; the characterization in the text is therefore a reconstruction from his position elsewhere, and a reader wanting the appraisal's own words will have to find the pamphlet.

[29]Michael E. Gordon, "An Evaluation of Jaques' Studies of Pay in the Light of Current Compensation Research," *Personnel Psychology* 22, no. 4 (1969): 369–89. The three routes are at 386–87: the behavior of pay over time, statistical models of pay, and the relationship between pay and other job-satisfaction variables. The concession is in Craddock's *Requisite Organization Annotated Bibliography*, in the entry for this article: "Gordon's other research suggestions have yet to be explored." Craddock adds "HINT, HINT," which we have not quoted but which is the tone of the concession.

[30]Paul S. Goodman, "An Empirical Examination of Elliott Jaques' Concept of Time Span," *Human Relations* 20, no. 2 (1967): 155–70. The movement's own bibliography disputes the finding on two grounds — that Goodman obtained strong correlations when managerial roles

The advocates themselves report the internal weakness, and it is the one that matters most here. In Richardson's data, time-span correlated with felt-fair pay at 0.86 and with actual freedom to act at 0.41.[31] The measure predicts what people think a role should pay considerably better than it predicts the discretion the role is named after. For a theory of remuneration that is a curiosity. For our purposes it comes close to disqualifying, because the quantity we would be borrowing is discretion, and discretion is the half the instrument tracks least well.

Jaques then left the refereed journals and published through his own imprint and a practitioner network, and the movement maintains its own annotated bibliography.[32] None of that makes the theory false. The bibliography is also more honest than a partisan compilation has any obligation to be: the 0.41 figure, the unexplored validation routes and the limits of Richardson's reporting are all in it, entered by advocates rather than extracted by critics. It does mean the citation record cannot be read like a normal literature. Inside a network that publishes and cites itself, a count of citations measures the network rather than the claim. We enter that caution with no comfort at all. It describes a position one recursion level away from our own.

For the use the VSM would make of the ruler, a further study of pay would settle nothing. What would settle it is a construct-validity study of the boundaries: does the stratum a role falls into by time-span predict anything about the role other than what people think it should earn? Gordon named three such routes in 1969 and nobody ran them. We did not run them either, and we had as much reason to want the answer as the requisite-organization community did.

Holling and panarchy — the thing Beer's diagram cannot show

What they found · Phase-specific adaptive cycles nested across scales, coupled by revolt and remember. **When** · 1973–2002. **Who cited whom** · This literature postdates Beer's first VSM statement; it is convergence, not prior art. **Evidence base** · A review finds substantial metaphorical and

alone were used, and that the low result was not replicated in six later studies — and records that Goodman himself stepped back from the article in a 1973 book review, Paul S. Goodman and Robert L. Miller, review of *Fair Pay and Work*, by Roy Richardson, *Administrative Science Quarterly* 18, no. 1 (1973): 135–37. The six studies it lists are the Berkeley dissertations of Kvålseth (1971) and Wirth (1976), three Canadian studies by Hansen (1973, 1974, 1975), and Vinton's dissertation (1995). The text reports what Goodman reported; it does not adjudicate the dispute.

[31] Craddock, *Requisite Organization Annotated Bibliography*, entry for Richardson 1969 and 1971: the correlation of time-span with felt-fair pay is +0.862 and with freedom to act +0.407. Both figures come from the movement's own annotated bibliography, which is the point being made about them: the weaker number is reported by the advocates, not extracted by critics.

[32] Craddock, *Requisite Organization Annotated Bibliography*. Jaques's later imprint was Cason Hall, publisher of *Requisite Organization*; the practitioner network is the Global Organization Design Society, which hosts the bibliography.

retrospective use. **What it costs us** · Recursion without phase cannot show what panarchy's unequal couplings are meant to show.

Holling's adaptive cycle is the one member of this group that comes after Beer rather than before him: resilience in 1973, the cycle itself in 1986, panarchy proper at the turn of the century. Systems pass through exploitation, conservation, release and reorganization. The cycles nest across scales that differ in speed and extent. And the levels are coupled in both directions. In *revolt*, a fast collapse at a small scale cascades upward into a slower and larger level. In *remember*, the slow large level supplies accumulated structure downward to reconstitute a small one after its collapse.[33]

The contrast with Beer is the sharpest in this chapter, and it concerns time. Beer's recursion is structurally invariant: the same five systems at every level, and the diagram of a level looks the same whatever phase it is in. Beer did make intervention depend on purpose, since "Viable systems of concentrated purpose will be closely-knit" and those of general purpose "more loosely coherent", and in *Brain of the Firm* he distinguished modes of behavior for a whole firm, sustained activity, growth, retrenchment and crisis among them.[34] He has no phase that differs between levels at the same moment, and no coupling that changes with phase. Panarchy's nesting is temporal. Each level sits somewhere in a cycle, different levels are in different phases at once, and the coupling between them changes with phase. A System One↗ in release does not need the System Three↗ that a System One in conservation needs, and imposing the second on the first is a recognizable way to kill a unit that is trying to reorganize. A VSM diagram cannot represent that situation, and every restructuring produces it.

The channels are the second contrast. Beer's inter-recursion couplings (the command axis with its resource bargain, the planning he called the "glue" of cohesion, the cross-recursion coordination of System Two and Three-Star, and the algedonic signal) are all regulatory relations in his sense of the word.[35] Neither revolt nor remember is a regulatory relation. Revolt is contagion; remember is inheritance. The algedonic signal is arguably a degenerate revolt, a pain report that bypasses levels, though it is a report rather than a cascade, and it is addressed rather than propagated. Nothing anywhere in the model corresponds to remember. We therefore have no account of how a destroyed unit is reconstituted out of the memory held

[33]C. S. Holling, "Resilience and Stability of Ecological Systems," *Annual Review of Ecology and Systematics* 4 (1973): 1–23; C. S. Holling, "The Resilience of Terrestrial Ecosystems: Local Surprise and Global Change," in *Sustainable Development of the Biosphere*, ed. William C. Clark and R. E. Munn (Cambridge University Press, 1986), 292–317, where the adaptive cycle first appears; Lance H. Gunderson and C. S. Holling, eds., *Panarchy: Understanding Transformations in Human and Natural Systems* (Island Press, 2002), and C. S. Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405, for revolt and remember.

[34]Beer, *Heart of Enterprise*, 158; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 234–35.

[35]Beer, *Heart of Enterprise*, 336 and 407; Beer, *Diagnosing the System*, 38–39 and 135.

by its container, and the case is far from exotic. It follows every closure, every failed subsidiary and every team that is stood down and later stood up again.

On evidence, the panarchy community has done something the cybernetics community has not: it has audited itself in public. Its founders describe the framework as "not a theory of what it is, but a metaphor of what might be".[36] In 2022 a systematic review by resilience insiders coded a sample of the literature. Roughly half of it used panarchy purely as a metaphor, the empirical applications are dominated by qualitative and retrospective studies, and many papers confirm the framework's expectations without testing alternatives. Their central admission, as they wrote it: "There remains a lack of testable mechanistic theories to explain the dynamics of panarchy and, consequently, there are a lack of methods to measure or observe it in real social-ecological systems."[37] Nearly every application is retrodiction.

The parts of the framework in actual use are also, for our purposes, the wrong parts. The adaptive cycle appears in about four fifths of the coded papers, the rigidity trap in a fifth, the poverty trap in a seventh. Cross-scale interaction, the nesting that makes panarchy *panarchy* rather than the adaptive cycle, ranks second, behind the cycle, and the review records that papers did not engage in measurement or observation to test revolt and remember.[38] The component we would most want to borrow is the one that community has tested least.

---

[36] Gunderson and Holling, *Panarchy*, 32.

[37] Juan C. Rocha, Linda B. Luvuno, Jesse T. Rieb, Erin T. H. Crockett, Katja Malmborg, Michael Schoon, and Garry D. Peterson, "Panarchy: Ripples of a Boundary Concept," *Ecology and Society* 27, no. 3 (2022): art. 21, https://doi.org/10.5751/ES-13374-270321. The coded sub-sample is N = 42: "Roughly half (20/42) of the papers used panarchy as a metaphor"; "Empirical papers predominantly used qualitative methods (17) and were primarily retrospective historical reviews"; "Many of the papers we read confirmed but did not challenge the expectations that panarchy put forward." A correction the reader is owed: earlier drafts of this chapter attributed to this paper the sentence "we still lack the theory and methods to gather observations and be able to decide in which phase of the adaptive cycle a system currently is." That sentence is not in the paper. We searched the full text of the article and of the one other work it could plausibly have come from, zu Castell and Schrenk's, and searched for the phrase in the open literature, and we could not source it anywhere. It has been replaced above by the review's own statement of the same admission, and the earlier sentence should not be quoted by anyone as Rocha and colleagues' words.

[38] This is a correction. Rocha and colleagues report the adaptive cycle at 81 per cent of the coded sample (34 of 42), the rigidity trap at 19 per cent and the poverty trap at 14 per cent, and they rank cross-scale interactions second: "the most common feature of panarchy in the literature was the adaptive cycle (81%, 34 out of N = 42) followed by cross-scale interactions." Earlier drafts of this chapter called the nesting the least-used component of all, which the review contradicts. The review's conclusion nonetheless supports the weaker claim now made — "Cross-scale dynamics, structure, and traps have received less attention" — and it records that papers "did not engage in measurement or observation to test… revolt and remember cross-scale interactions." The review's figure 5, which breaks the coding down further, has no separate code for nestedness: hierarchy and cross-scale interactions are coded together, at about three quarters of the sub-sample, second only to the adaptive cycle.

There is also one real test, and it went against the framework. In 2020 zu Castell and Schrenk built an information-theoretic method for identifying the phases and applied it to the European financial crisis and to grassland ecosystems. They recovered the phases, and reported: "We cannot confirm a generally short duration of the Ω- and α-phase compared to the r- and K-phase."[39] The fast back loop and slow front loop drawn into every version of the diagram did not appear in the data. The authors also warn that their method cannot be used in a black box fashion and depends on the analyst's choice of parameters. That is Simon's threshold problem, arriving at the other end of the chapter under a different name.

The VSM's own epistemic position is very nearly the same: a rich generative model, wide uptake, thin falsification. The difference is the one this book is about. Somebody built an instrument, the instrument returned a negative result, and the community published it.

What the four have in common

Taken together, the four entries yield a single result, and it has nothing to do with any one author's carelessness.

Simon has a threshold nobody can set, and says so. Koestler has propositions nobody can test, and does not say so. Jaques has a ruler whose construct validity his own side concedes is weak, and records the concession in his own bibliography. Holling has a phase nobody has yet shown they can identify in the present tense, and the review that establishes this was written by insiders. Beer has a criterion that resolves to a judgment call, which is the finding of Chapter 5.

Four people saw related structures independently, and Beer saw his own version a fifth time. None of the five could measure it. The instrument has been missing for six decades in five literatures at once, in fields with different data, different methods, different funding and different standards of proof. That distribution is itself evidence, and it points away from the explanation each community reaches for first. The difficulty is unlikely to lie in cybernetics, or in requisite organization, or in resilience ecology. It lies in the structure, and whoever builds the instrument will have done something considerable for all five.

[39]Wolfgang zu Castell and Hannah Schrenk, "Computing the Adaptive Cycle," *Scientific Reports* 10 (2020): art. 18175, https://doi.org/10.1038/s41598-020-74888-y. The method is information-theoretic, built on transfer entropy in a network representation, and is applied to a Tangled Nature model, to the European financial crisis and to an ecological succession experiment in Kansas grassland. The quotation is verified verbatim, and its context is worth keeping: "Some of the details that Gunderson and Holling draw on cannot be supported by our study. For example, we cannot confirm a generally short duration of the Ω- and α-phase compared to the r- and K-phase." The authors also warn that the method "can not be used in a black box fashion" and depends on the analyst's choice of parameters — which is Simon's epsilon problem, arriving at the other end of the chapter.

Two readings of the situation hold up, and Beer held both. In the first, levels are real and hard to measure, as many real things are, and the convergence of five independent traditions on related objects that none of them can measure is a reason to keep working. In the second, levels are an artifact of how observers parse continuous systems: five traditions converged because five sets of humans parse in the same way, and the instrument is missing because there is nothing for it to measure. Beer wrote the second as method: "A system is not something given in nature, but something defined by intelligence"; two people "will not necessarily agree on the existence, or nature, or boundaries of any system so detected"; and it is possible to say only that treating a collection of things as a system "is helpful." He wrote the first when he said that levels of recursion "answer to mathematical criteria."[40] Neither half has been put to a test. We believe the first. We cannot presently show that the second is wrong, and no one in the other four literatures can either. That is why the instrument matters more than another round of exposition.

This part of the book is no priority dispute, and the four neighbors do not compete for our territory. Each offers the model something specific that it lacks. Simon offers a boundary criterion stated over observable rates rather than over identity. Koestler offers a noun for the recursing thing and a set of channel rules stated sharply enough to be broken. Jaques shows that a level, in his sense of the word, can be attached to a number at all, and leaves an honest record of what that number does and does not predict. Holling offers phase, and the two cross-scale couplings we have no counterpart for. We did not read this literature, and it cost us four opportunities to test our own central construct against somebody else's.

The priced receipt

Receipt — one registered blind comparison of level boundaries, in one organization. Build an interaction matrix from observed rates of interaction. Record time-span of discretion for the roles that head each candidate recursion's management, at the boundaries an experienced VSM diagnostician independently selects. As a third and cheap boundary reading, have the heads of each candidate recursion complete Beer's identity exercise at the same boundaries, which puts Beer's own nearest thing to a ruler beside time-span; Beer called what it asks for "subjective judgments."[41] Then give the matrix and the VSM account to two analysts who do not know which method produced which boundaries, and ask each to reconstruct the level structure. Register the protocol and the criterion of agreement before any of it is run. Agreement is measured as the adjusted Rand index between the

[40] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 242 and 243, paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994); Beer, *Heart of Enterprise*, 308.

[41] Beer, *Heart of Enterprise*, 417–33, "Identity: an exercise"; the caution is at 419: "These are subjective judgments for which we ask."

matrix-derived and the VSM-derived partitions of units; chance is estimated by a permutation test against 1,000 random partitions with the same block sizes, and convergence requires p below 0.05 and an index of at least 0.5, both registered before the run. Ten interview-days, five observation-days and five analysis-days make this approximately twenty working days.

The design stands or falls on a negative result being publishable in advance. If the boundaries do not converge beyond chance, the claim that recursion identifies an observable level structure is weakened, and so is the assumption that Jaques's strata and Beer's recursions can be read off the same boundaries. We should say so in print. If they do converge, the field has the beginning of a ruler, and the first evidence in sixty years that the thing five literatures have been pointing at can be seen twice, by two methods.

Chapter 12

The Variety Logic: Galbraith, Boisot, Aulin, Friston

In 1968 Jay Galbraith spends a little over two months in residence at the Commercial Airplane Division of the Boeing Company in Seattle, with two graduate students. They interview, read internal documents and sit in meetings. Boeing management reads the text before it goes out as a Sloan School working paper.

He takes away two things. The structure of an organization is set by how much information it has to process under uncertainty. And a designer has two families of move and no others: reduce the need for information, or increase the capacity to handle it.

Those are variety↗[1] attenuation↗ and variety amplification. The paper does not contain the word *variety*. It does not contain *Ashby*, *Wiener*, *Shannon* or *cybernetic* either, and neither does the article that followed six years later. We searched both in full.

He got there from a factory in Seattle, and he got there without us.[2]

The oldest borrowing

The variety logic is the oldest thing in this book and the least Beer's own, and he never pretended otherwise. Ashby is in every one of his major works, and requisite variety↗ is the load the whole structure rests on. Chapter 10 found him among the Explicit References of *Brain of the Firm* and moved on.

Chapter 3 took the narrower question: what Ashby proved, what Beer's *absorb* compresses, and how little use we made of the one adjacent law that

[1]Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

[2]Jay R. Galbraith, *Environmental and Technological Determinants of Organization Design: A Case Study*, Working Paper 352-68 (Sloan School of Management, MIT, 1968), https://dspace.mit.edu/handle/1721.1/47606, for the residence at Seattle, the method and the review by Boeing management; Jay R. Galbraith, "Organization Design: An Information Processing View," *Interfaces* 4, no. 3 (1974): 28–36, for the two families of move. The negative search covers the working paper and the article, both read in full; for *Designing Complex Organizations* (1973), whose reference list we could not obtain in a readable form, it rests on a search of a full-text index only. The scope of that finding is set out below.

would have licensed the recursion we teach. None of that is restated here. The fact at issue now is harder to hold in one piece. The most striking evidence that the variety logic is real is that a management theorist rederived it in the early seventies from field data, having apparently never read a word of cybernetics. That same evidence shows that Beer's vocabulary was sufficient to reach the logic and never necessary to it.

Our teaching tends to split that sentence in two. We quote the rediscovery when we want to show that the principle is no private notation, and we set it aside when we want to claim the principle as ours. This part of the book asks us to hold both halves at once. A literature that arrived at the same design space without our vocabulary is a corroboration and a limit in the same document, and the limit is the part we have never priced.

Each case below goes on the same ledger: what was found, when, who cited whom, what the evidence base actually was, and what the finding costs the cybernetic claim. The last column is the one we have been slowest to fill in.

Galbraith: the same design space, in another vocabulary

Galbraith published *Designing Complex Organizations* in 1973, and its compressed statement, "Organization Design: An Information Processing View", in *Interfaces* the following year.[3] The premise is a proportionality: "the greater the uncertainty of the task, the greater the amount of information that has to be processed between decision makers during the execution of the task."[4] The design consequence follows directly. An organization has two ways to reduce the information it must process (slack resources, self-contained tasks) and two ways to raise the capacity with which it handles what remains (vertical information systems, lateral relations).

That is attenuation and amplification with the labels changed, and the correspondence is close enough to state carefully. Galbraith's variable is uncertainty. Beer did treat variety, in bits, as uncertainty; what neither he nor we have supplied is the argument that Galbraith's task uncertainty is that quantity.[5] The shape of the design space he describes is unambiguous: two directions of move, claimed as exhaustive, distinguished by which side of a mismatch they act on. A designer who has never encountered the word *variety* is working the same two levers, and knows there are only two. His

[3] Jay R. Galbraith, *Designing Complex Organizations* (Addison-Wesley, 1973); Jay R. Galbraith, "Organization Design: An Information Processing View," *Interfaces* 4, no. 3 (1974): 28–36, https://doi.org/10.1287/inte.4.3.28.

[4] Galbraith, "Organization Design," 28. The four strategies — creation of slack resources, creation of self-contained tasks, investment in vertical information systems, creation of lateral relationships — are set out at 29–35 and are claimed at 36 to be "an exhaustive set of alternatives."

[5] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 252, for the logarithmic measure of variety, and 492–93, where the entropy of a decision, in bits, is "the measure of uncertainty" and "43 bits is a measure of variety"; paginated after the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994).

two capacity moves, vertical information systems and lateral relations, are closer to Beer's Second Principle, that channels must carry more variety than their source generates, than to amplification.[6]

On one point his formulation is sharper than the one we usually quote. "The task information requirements and the capacity of the organization to process information are always matched. If the organization does not consciously match them, reduced performance through budget overruns, schedule overruns will occur in order to bring about equality."[7] Ashby's law says that regulation fails when variety is short. Galbraith says that the books always balance, and names the currency in which the balancing is paid: schedule, budget, performance. Beer had said as much about the balance. His First Principle of Organization holds that managerial, operational and environmental varieties "tend to equate; they should be designed to do so with minimal damage to people and to cost."[8] Galbraith's advance is the currency. Beer names damage to people and to cost in general terms; Galbraith names overruns an observer can check, and with them the claim takes an empirical shape. It tells an observer where to look when a design is under-specified, and what to expect there. Whether it is true can be put to data. Our own statements of the same idea have rarely been cast in a form that would let anyone ask.

This settles less than it might seem to. The paper asserts the matching claim without demonstrating it, and the currencies it names are those of an aerospace program in the late sixties, which is where Galbraith had been standing. The formulation does not prove the variety logic. Its value lies in stating the logic in terms a manager can act on and a researcher can measure, which is the step we did not take.

The lineage we can verify, and the part we cannot

The claim that Galbraith arrived without cybernetics is negative, and a negative claim is only as good as the scope attached to it. The *Interfaces* paper carries nine references: Cyert and March, Lawrence and Lorsch, Lorsch and Lawrence, March and Simon, Pounds, Simon, and three items of his own. The words *Ashby*, *Wiener*, *Shannon*, *cybernetics* and *variety* do not occur in it at all; his word throughout is *uncertainty*. The 1968 working paper behind the book is the same. We searched both in full.[9] He states

[6] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 99, the Second Principle of Organization: the channels "must each have a higher capacity to transmit a given amount of information relevant to variety selection in a given time than the originating sub-system has to generate it in that time." The Third Principle, on transduction, follows at 101.

[7] Galbraith, "Organization Design," 36.

[8] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 97; the glossary of rules at 565 prints "minimum damage".

[9] The negative finding is direct for the 1974 *Interfaces* article and for the 1968 MIT working paper that preceded the book, both of which we searched in full: neither contains the strings *Ashby*, *Wiener*, *Shannon*, *cybernetic* or *variety*. The nine items in the *Interfaces* bibliography (36) are Cyert and March, *The Behavioral Theory of the Firm* (1963), so titled by Galbraith though

his lineage explicitly, and it points somewhere else entirely: "the cognitive limits theory of Herbert Simon was the guiding influence."[10] That is the Carnegie behavioral tradition rather than the Macy one, a lineage with its own account of why organizations take the shapes they do.

The scope has a limit, and we would rather state it than have it found. We have not been able to read the reference list of *Designing Complex Organizations*. A search of the full-text index of the Internet Archive's scan of the book finds neither Ashby nor Beer, Wiener, Shannon or *cybernetic*, but a search of an index is weaker than a reading. The negative finding therefore holds for the 1974 article and the 1968 working paper, and holds provisionally for the 1973 book. Anyone repeating the claim should repeat it at that scope, and the sentence in this chapter is written so that it can be.

A second limit is more general, and it survives even if the book turns out to be clean. Absence of citation does not establish absence of influence. By the late sixties the vocabulary of information, channel and capacity was in wide circulation, and a researcher could take up the frame without ever meeting its sources. The negative search establishes that Galbraith did not derive his argument from ours and did not need ours to make it stand. The stronger question would be settled by documents rather than by texts: correspondence, teaching materials, the reference list we could not obtain. Until someone does that work, the defensible formulation is independence of derivation, not innocence of the surrounding air.

The evidence base needs the same care. It was a longitudinal case study inside a single division: Boeing's commercial airplane division in Seattle, two months in residence with two graduate students, interviews, internal documents, observation of meetings, and the text reviewed by management

---

the book is *A Behavioral Theory of the Firm*; Galbraith, "Environmental and Technological Determinants of Organization Design: A Case Study," in Lawrence and Lorsch, ed., *Studies in Organization Design* (1970); Galbraith, "Designing Matrix Organizations" (1971); Galbraith, *Organization Design* (1973), his own short title for *Designing Complex Organizations*; Lawrence and Lorsch, *Organization and Environment* (1967); Lorsch and Lawrence, "Environmental Factors and Organization Integration," a 1968 conference paper; March and Simon, *Organizations* (1958); William Pounds, "The Scheduling Environment," in John F. Muth and Gerald L. Thompson, eds., *Industrial Scheduling* (Prentice-Hall, 1963); and Simon, *Models of Man* (1957). The Pounds item is the scheduling chapter, not his better-known "The Process of Problem Finding"; Galbraith cites it once, at 30, for job-shop delivery quoting. The reference list of *Designing Complex Organizations* (1973) we were unable to obtain in a form we could read. A search of the full-text index of the Internet Archive's scan of the book (item designingcomplex00galb, run on 29 September 2026) returns no hit for *Ashby*, *Beer*, *Stafford*, *Wiener*, *Shannon*, *cybernetic*, *cybernetics* or *requisite*, and hits for *Simon*, *Lawrence*, *Thompson* and *Boeing*; the index cannot show where a word occurs or how the scan was read. The claim in this chapter is therefore that Galbraith's cybernetic silence is verified for the paper and the working paper and provisional for the book.

[10] Galbraith, "Organization Design," 36. In full: "The purpose of this paper has been to explain why task uncertainty is related to organizational form. In so doing the cognitive limits theory of Herbert Simon was the guiding influence." The Simon in the bibliography is Herbert A. Simon, *Models of Man: Social and Rational* (John Wiley & Sons, 1957), together with James G. March and Herbert A. Simon, *Organizations* (John Wiley & Sons, 1958).

before release.[11] One site, one program, one period. The paper's own contrast is with the cross-sectional data of Lawrence and Lorsch, and the single case bought depth at the cost of range. Stating this keeps us from overclaiming on behalf of a literature that is not ours to overclaim on. It also makes the comparison with our own record exact. A two-month case study inside one division is a modest evidentiary base, and it is still more field data than most of our own claims have ever been asked to rest on.

What the rediscovery costs, and what it does not

Two conclusions follow, and only one of them is comfortable.

The comfortable one is that the variety logic is not arbitrary. It was assembled from parts that other literatures found independently, so nobody invented it as a notation and then persuaded people to adopt it. When the model is charged with arbitrariness, with components chosen because Beer chose them, the independent arrival at the same two-move design space is the strongest answer available to us. It is a good answer, and we should keep giving it.

The uncomfortable one is the bill attached to the first. The fact that clears us of arbitrariness also denies us proprietorship. If the logic can be reached from a factory in Seattle by way of Herbert Simon, then the problem imposes it on anyone who looks hard enough at how organizations cope under uncertainty. Our claim on it is one of priority and of completeness: that we arrived first, and that we hold the logic in a more general form. Neither amounts to ownership. Priority is a historical question, and it settles largely in Ashby's favor rather than ours. Completeness is a question about what the model adds once the variety logic has been granted, and it belongs to Part Four, where it has to be answered narrowly or not at all.

Two neighboring derivations

Two further arrivals belong on the same ledger, for different reasons.

In 1999 Boisot and Child distinguished complexity reduction from complexity absorption as alternative adaptive strategies: a system can cut the complexity it confronts, or it can take that complexity in and hold it in a

[11] Jay R. Galbraith, *Environmental and Technological Determinants of Organization Design: A Case Study*, Working Paper 352-68 (Sloan School of Management, MIT, 1968), https://dspace.mit.edu/handle/1721.1/47606. Abstract: "This paper reports a longitudinal case study of the Commercial Airplane Division of the Boeing Company." The opening note records that the study was done with the cooperation of that division at Seattle, that Galbraith and two graduate students spent "a little over two months in residence," and that the data were interviews, internal Boeing documents and observation of meetings, with the text reviewed by Boeing management. This is a single-site longitudinal case, not a multi-firm survey, and the divisional restriction matters: the paper's contrast is with the cross-sectional data of Lawrence and Lorsch.

repertoire of responses.[12] Theirs is a neighboring literature with its own lineage, the Information-Space framework, and nothing in the paper suggests a dependency on Beer. Fairness to them as much as to us requires an accurate account of the evidence base. It is a conceptual paper illustrated at length with China rather than an analysis of a Chinese institutional dataset, and it is stronger on drawing distinctions than on measurement. The cost to us is modest but real. The attenuation–amplification distinction travels perfectly well without the VSM attached to it, in a vocabulary organization theorists already had to hand.

Aulin-Ahmavaara's law of requisite hierarchy appeared in *Kybernetes* in 1979, the year of *The Heart of Enterprise*.[13] It holds that hierarchy compensates when regulatory ability is limited, which reads in outline like a derivation of what Beer asserted. Chapter 3 records why that reading may be overstated: the hierarchy Aulin-Ahmavaara derives is one of control, and the paper mentions neither Beer nor recursion. It sits inside cybernetics, so it is no rediscovery in the Galbraith sense, and this chapter establishes no derivation in either direction between it and Beer's recursion. Chapter 3 has already recorded how little the law was taken up. The figure appears here only to fix an order of magnitude: 44 citations recorded by OpenAlex on 29 September 2026, from a single database, an instrument crude enough that we use it for nothing finer.[14] For this chapter the lesson is narrow. A neighboring formal result existed for a claim related to the recursion we teach and not identical with it, since Beer's recursion rests on viable systems inside viable systems rather than on limited regulatory ability.[15] It was published in a cybernetics journal in the same year as *The Heart of Enterprise*, and first stated four years earlier, and we did not build on it.

[12]Max Boisot and John Child, "Organizations as Adaptive Systems in Complex Environments: The Case of China," *Organization Science* 10, no. 3 (1999): 237–52, https://doi.org/10.1287/orsc.10.3.237 — a conceptual paper built on their Information-Space framework and illustrated with China at length, rather than an analysis of a Chinese dataset; A. Y. Aulin-Ahmavaara, "The Law of Requisite Hierarchy," *Kybernetes* 8, no. 4 (1979): 259–66, https://doi.org/10.1108/eb005528; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979).

[13]Max Boisot and John Child, "Organizations as Adaptive Systems in Complex Environments: The Case of China," *Organization Science* 10, no. 3 (1999): 237–52, https://doi.org/10.1287/orsc.10.3.237 — a conceptual paper built on their Information-Space framework and illustrated with China at length, rather than an analysis of a Chinese dataset; A. Y. Aulin-Ahmavaara, "The Law of Requisite Hierarchy," *Kybernetes* 8, no. 4 (1979): 259–66, https://doi.org/10.1108/eb005528; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979).

[14]On 29 September 2026, OpenAlex recorded 44 citations of Aulin-Ahmavaara's paper, against 1,065 for Conant and Ashby (1970) and 2,879 for Francis and Wonham (1976); on 7 August 2026 the counts were 43, 1,040 and 2,838. Citation counts from one database are a crude instrument and we use them only for the order of magnitude.

[15]Beer's criterion is stated in Chapter 10 (Stafford Beer, *Diagnosing the System for Organizations* [John Wiley & Sons, 1985], 2; Beer, *Heart of Enterprise*, 118 and 308).

Friston: the mathematics Beer never had

The free energy principle is the first attempt since Ashby to derive regulatory structure from first principles about boundary maintenance with mathematics that runs. Beer's derivation in *The Heart of Enterprise* had the same aim and ended in balance equations and a theorem. Take any system that persists and whose states divide by a Markov blanket into internal, sensory, active and external. It will appear on average to minimize variational free energy, an upper bound on the surprisal of its sensory states, so that its internal dynamics can be read as approximate Bayesian inference about its environment.[16] Active inference extends this to action: the system either changes its beliefs to match the world or changes the world to match its beliefs.[17]

It is tempting to call this variety attenuation and amplification restated with better mathematics. That is not quite right. The belief/action duality differs from the attenuate/amplify one. The two can be brought into correspondence, but doing so is an interpretive move, and it should be labeled as one rather than presented as a derivation. What they genuinely share is the ambition: a claim about any boundary-maintaining system, at any scale, in any substrate. That is Beer's ambition, arriving a generation later with an apparatus he did not have.

Two things follow from that apparatus, and they point in opposite directions.

The first is that it runs. There is working code, there are published simulations, and there is a pipeline that leads from a generative model through simulation to model inversion against behavioral data.[18] We have not run any of it ourselves and make no claim about the quality of any particular fit; we claim only that the pipeline exists. A VSM↗ diagram cannot be executed. Nobody has ever shown that a VSM claim is mathematically wrong, and survival under scrutiny of that kind has nothing to do with it. Beer wrote

[16] Karl Friston, "The Free-Energy Principle: A Unified Brain Theory?," *Nature Reviews Neuroscience* 11, no. 2 (2010): 127–38, https://doi.org/10.1038/nrn2787; Karl Friston, "Life as We Know It," *Journal of the Royal Society Interface* 10, no. 86 (2013): 20130475, https://doi.org/10.1098/rsif.2013.0475; Karl Friston, "A Free Energy Principle for a Particular Physics," arXiv preprint arXiv:1906.10184 (2019), https://doi.org/10.48550/arXiv.1906.10184. The Markov-blanket statement paraphrased in the text is the 2013 paper's: the blanket "separates internal and external states in a statistical sense," internal states "will appear to minimize a free energy functional of the states of their Markov blanket," and because this is the quantity optimized in Bayesian inference they "will appear to engage in active Bayesian inference." The word doing the work in all three papers is *appear*.

[17] Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo, "Active Inference: A Process Theory," *Neural Computation* 29, no. 1 (2017): 1–49, https://doi.org/10.1162/NECO_a_00912; Thomas Parr, Giovanni Pezzulo, and Karl J. Friston, *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (MIT Press, 2022), https://doi.org/10.7551/mitpress/12441.001.0001.

[18] Friston et al., "Active Inference: A Process Theory"; Parr, Pezzulo, and Friston, *Active Inference*. Our claim is about the existence of an executable pipeline — generative model, simulation, model inversion against behavioral data — as reported in these sources. We have not run the code ourselves and make no claim about the quality of any particular fit.

his central claims as equations and a theorem, and of the theorem he wrote: "The proof offered later will be a topological proof." He also wrote that the equalities "require some *extra* variety, which is called redundancy", and he did not measure the terms.[19] The claims look precise and cannot be checked, which is a different condition from being wrong. That asymmetry matters more than any particular result in the free energy literature.

Explicit enough to be wrong

The second is that being explicit has a price, which is being corrected in public. The free energy literature has paid it, and ours has not.

In 2021 Biehl, Pollock and Kanai showed that the definitions of the Markov blanket are not equivalent across the corpus, and proved by counterexample that the original free energy lemma, taken at face value, is incorrect. They were explicit that their target was the formal argument rather than the underlying ideas.[20] The following year Aguilera, Millidge, Tschantz and Buckley showed that even for weakly coupled linear systems the conditions the argument requires hold only within a very narrow band of parameters.[21] Bruineberg and colleagues, in *Behavioral and Brain Sciences*, separated "Pearl blankets" from "Friston blankets". The first is the Markov blanket as Judea Pearl defined it in 1988, an epistemic tool for inference in a network; the second is a metaphysical construct used to demarcate where an agent ends and the world begins. They argued that the literature persistently trades on the conflation.[22] And Friston himself, asked directly, has said of the principle, as quoted in Chapter 9: "Like Hamilton's Principle of Stationary Action, it cannot be falsified. It cannot be disproven."[23]

---

[19] Beer, *Heart of Enterprise*, 118 for the theorem and the announced proof; 217 and 298 for the axioms; 355 for the redundancy qualification. A full-text search of the book (29 September 2026) finds no passage labeled as that proof; the topological constructions at 289 and 346 may be what Beer had in mind, and the text of this chapter claims only that the proof was announced and that the terms were not measured.

[20] Martin Biehl, Felix A. Pollock, and Ryota Kanai, "A Technical Critique of Some Parts of the Free Energy Principle," *Entropy* 23, no. 3 (2021): 293, https://doi.org/10.3390/e23030293. They report that "various definitions of the 'Markov blanket' proposed in different works are not equivalent" and that "the original free energy lemma, when taken at face value, is wrong," and they are explicit that the target is the formal argument, not the underlying ideas.

[21] Miguel Aguilera, Beren Millidge, Alexander Tschantz, and Christopher L. Buckley, "How Particular Is the Physics of the Free Energy Principle?," *Physics of Life Reviews* 40 (2022): 24–50, https://doi.org/10.1016/j.plrev.2021.11.001. The analysis is of weakly coupled linear (Ornstein–Uhlenbeck) systems, and the conclusion is that the Markov blanket condition together with the restrictions on solenoidal flow "are only valid for a very narrow space of parameters."

[22] Jelle Bruineberg, Krzysztof Dołęga, Joe Dewhurst, and Manuel Baltieri, "The Emperor's New Markov Blankets," *Behavioral and Brain Sciences* 45 (2022): e183, https://doi.org/10.1017/S0140525X21002351.

[23] Karl Friston, Martin Fortier, and Daniel A. Friedman, "Of Woodlice and Men: A Bayesian Account of Cognition, Life and Consciousness; An Interview with Karl Friston," *ALIUS Bulletin* 2 (2018): 17–43, at 21. The context is his distinction between a normative principle and a process theory: "the free energy principle is what it is—a principle. Like Hamilton's Principle of Stationary Action, it cannot be falsified. It cannot be disproven. In fact, there's not much

We enter none of this as a defeat for anyone. It describes a working relationship between a general claim and the criticism such a claim can attract, and the central concession is available in the author's own words. The framework carries Beer's ambition with mathematics Beer never had, and its architect declines to call it falsifiable. What survives the criticism is a generative machinery for building testable process theories. The claim to be a law does not survive. We recognize that position because it is our own, except that our version was written as equations and never in terms anyone could check, so it did not attract the criticism in the first place. Being told that your lemma is wrong is a cost. It is also a form of participation, and by staying imprecise we have avoided paying for it.

The blanket is not ours

One genealogical point needs recording, because claiming neighboring work as cybernetic inheritance is a habit we should not indulge. Markov blankets are Judea Pearl's, from 1988, and their route into active inference runs through Bayesian networks and variational inference.[24] Cybernetics did not bequeath them. The free energy literature does cite Conant and Ashby, but for the model claim rather than for the blanket, and running the two together produces a genealogy that will not hold.

The instrument we did not build

Galbraith's descendants spent four decades operationalizing information-processing requirements, information-processing capacity and the fit between the two, with validated constructs, in a line of work running from Tushman and Nadler in the late seventies to supply-chain studies four decades later.[25] Those are three points on a

you can do with it, unless you ask whether measurable systems conform to the principle." He goes on to argue that demanding falsifiability of a principle is a category error.

[24]Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, 1988). The tracing of the construct from Bayesian networks through variational inference to active inference is Bruineberg et al., "The Emperor's New Markov Blankets." That the free energy literature cites Conant and Ashby for the model claim rather than for the blanket can be checked in Friston, "A Free Energy Principle for a Particular Physics," which invokes "the good regulator theorem (Conant and Ashby, 1970)" alongside the Bayesian brain hypothesis and nowhere for the blanket.

[25]Michael L. Tushman and David A. Nadler, "Information Processing as an Integrating Concept in Organizational Design," *Academy of Management Review* 3, no. 3 (1978): 613–24, https://doi.org/10.5465/amr.1978.4305791; G. Premkumar, K. Ramamurthy, and Carol Stoak Saunders, "Information Processing View of Organizations: An Exploratory Examination of Fit in the Context of Interorganizational Relationships," *Journal of Management Information Systems* 22, no. 1 (2005): 257–94, https://doi.org/10.1080/07421222.2003.11045841 (the DOI stem carries 2003; the article appeared in 2005); Ravi Srinivasan and Morgan Swink, "An Investigation of Visibility and Flexibility as Complements to Supply Chain Analytics: An Organizational Information Processing Theory Perspective," *Production and Operations Management* 27, no. 10 (2018): 1849–67, https://doi.org/10.1111/poms.12746. These are three points on a line, not a survey of it.

line rather than a survey of it, and the line is not uniformly strong. The direction of travel is unambiguous all the same, and we did not travel it. Nothing in principle prevents the variety claims from being given the same treatment. Disturbance variety, attenuators, amplifiers and response capacity are candidate constructs, and construct validation is a well-understood craft with a standard apparatus. After fifty years, the subject matter cannot be blamed for the absence of that work. We chose it by default, and the remedy is a study rather than an argument.

Beer anticipated part of the difficulty. Of his own variety calculations he asked: "Could we ever say: *this* is the variety of the chemical industry, say? Of course we could not." The use of the measure, he went on, "is not the naive use of counting states, but of matching state generators."[26] A negative result from the study priced below would refute a stronger claim than the one Beer made, that variety is one measurable organizational quantity. The balance claim would then have to be stated in terms of gaps an observer can see, and nobody has yet done that.

The priced receipt

Receipt — construct validation of the variety quantities. Adapt established organizational-information-processing scales[27] into a construct-validation study of disturbance variety, attenuators, amplifiers and response capacity. Design: a pilot in three organizations, comprising instrument design, thirty cognitive interviews, and two survey waves. Price: roughly forty researcher-days plus participant time. Pre-register convergence and discriminant-validity thresholds before the first wave. Negative result: if the proposed variety measures do not converge, stop treating "variety" as one measurable organizational quantity and specify the narrower constructs that do.

[26] Beer, *Heart of Enterprise*, 86–87.

[27] For example the measures of information-processing needs, capability and fit in Premkumar, Ramamurthy and Saunders, "Information Processing View of Organizations" (2005), cited in the note on Galbraith's descendants above.

Chapter 13

The Autonomy Design: Trist, Herbst, Cherns, Boeke, Endenburg

Ken Bamforth, a former miner then studying at the Tavistock Institute, goes back down a mine. At the Haighmoor seam he finds something the industry's own engineers had not designed: a set of relatively autonomous groups, interchanging roles and shifts, regulating their affairs with a minimum of supervision.

He brings it to Eric Trist. What they publish in 1951 claims more than that teams are pleasant: for a given technology there is a genuine organizational *choice*, and the design which jointly optimizes the social and the technical beats the one that optimizes the technical alone.

Trist, writing it up thirty years later, is exact about what was seen: "Co-operation between task groups was everywhere in evidence; personal commitment was obvious, absenteeism low, accidents infrequent, productivity high."

That is what an observation looks like when somebody has gone and made it. It is 1951. *Cybernetics and Management* is eight years away and *Brain of the Firm* twenty-one.[1]

The dial that has to be set

The autonomy-cohesion problem↗[2] is the one Beer addressed most convincingly, and also the one on which the literatures outside cybernetics traveled furthest without him.

The problem is easy to state and hard to escape. Give the parts of a system enough freedom to absorb their own variety↗ and they drift out of alignment with each other and with the whole. Take the freedom back and

[1] E. L. Trist and K. W. Bamforth, "Some Social and Psychological Consequences of the Longwall Method of Coal-Getting," *Human Relations* 4, no. 1 (1951): 3–38; Eric Trist, *The Evolution of Socio-Technical Systems*, Occasional Paper No. 2 (Ontario Quality of Working Life Centre, 1981), 8, for both quoted passages — the description of the Haighmoor groups and the sentence on what was observed. The dates measured against are Stafford Beer, *Cybernetics and Management* (English Universities Press, 1959), and the first edition of *Brain of the Firm* (Allen Lane, 1972).

[2] Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

the center has to absorb variety it cannot possibly hold. Everything Beer built between System One↗ and System Three↗ is an attempt to hold that dial in a defensible position: the principle of autonomy, the anti-oscillation function, the resource bargain, the audit channel that reaches past the manager.

On the requirement, that apparatus is strong. It treats autonomy as a quantity to be allocated, not a virtue to be conferred. The allocation is paid for somewhere, and the bill arrives either as variety the center cannot process or as oscillation between the parts, depending on which way the setting is wrong. Very little management writing gets that far, and it is a large part of why we keep teaching the model.

On the setting itself Beer says more than we teach and less than a designer needs. He claims the setting is determinate: "*Freedom is in principle a computable function of systemic purpose as perceived*," and the metasystem should make "only that degree of intervention that is required to maintain cohesiveness."[3] He assigns the choosing of it to the resource bargain, the "deal" by which "some degree of autonomy is agreed" between senior management and its junior counterparts.[4] He drafted a questionnaire that asks managers to rate how far the whole and its elements agree on the degree of autonomy, and hoped that "maybe (one day) thresholds will be established."[5] He introduced it with a warning of his own: "it would be bogus to pretend that proper measurements can be made. These are subjective judgments for which we ask."[6] He did not supply the computation, a validated measure or the thresholds, and nobody in the tradition has supplied them since. Chapter 4, *Channels Without Owners*, made the Part One version of this charge about the channels themselves. The same gap runs through the autonomy question, and restating the principle more emphatically does not close it. We can say that the metasystem, in practice mainly System Three, should spend no more of System One's autonomy than cohesion requires. We cannot say how much that is in a given case, we have no instrument that reports it, and we have not built one.

Three other traditions worked on the same dial. One of them produced a design rule two decades before *Brain of the Firm*. One produced a governance mechanism that a national labor authority will accept as a substitute for a works council. The third assembled an evidence base of a kind cybernetics has never assembled, and then spent its last decade warning people not to use it as a blueprint. That one is Elinor Ostrom's, and Chapter 14 takes it; this chapter takes the first two.

None of these traditions set out to correct us, and two of them appear not to have known we were there. Approached from a coal face, from a design office and from a family electrical engineering firm, the same dial

[3] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 158; restated at 200.

[4] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 38–39.

[5] Beer, *Heart of Enterprise*, 417–33 (the "Degree of Autonomy" questions, 424–25), with the sentence quoted at 432.

[6] Beer, *Heart of Enterprise*, 419.

produced components that fit inside our diagram without any help from us, and we did not read them. The loss ran both ways. They worked without the variety vocabulary that would have let them say what they had found; we worked without the field results and the design rules that would have let us specify what we were recommending.

Trist and Bamforth: the field finding

What they found · A semi-autonomous group inside a jointly designed technical and social whole. **When** · Haighmoor, 1951. **Who cited whom** · The lineage predates *Brain of the Firm*. **Evidence base** · Field observation at one coal seam and a long sociotechnical program. **What it costs us** · Autonomy becomes a design variable, not merely a permission in a diagram.

Eric Trist and Ken Bamforth published "Some Social and Psychological Consequences of the Longwall Method of Coal-Getting" in *Human Relations* in 1951, twenty-one years before *Brain of the Firm* and eight before *Cybernetics and Management*.[7] The subtitle does the theoretical work. The paper is about a work group "in Relation to the Social Structure and Technological Content of the Work System", which announces joint determination in the title. What Bamforth had seen at Haighmoor, groups interchanging roles and shifts and regulating their own affairs with almost no supervision, is what Trist summarized thirty years later: "cooperation between task groups was everywhere in evidence; personal commitment was obvious, absenteeism low, accidents infrequent, productivity high."[8] The Durham comparisons that most people mean when they say "the coal studies" came a decade later, in *Organizational Choice*.[9]

The claim, as the opening of this chapter gave it, is about organizational *choice*. Put in our terms, it says that variety absorption and the machinery

[7] E. L. Trist and K. W. Bamforth, "Some Social and Psychological Consequences of the Longwall Method of Coal-Getting: An Examination of the Psychological Situation and Defences of a Work Group in Relation to the Social Structure and Technological Content of the Work System," *Human Relations* 4, no. 1 (1951): 3–38. The dates against which we measure it are Stafford Beer, *Cybernetics and Management* (English Universities Press, 1959), and the first edition of Stafford Beer, *Brain of the Firm* (Allen Lane, 1972); the edition we cite elsewhere in this book is the second, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981).

[8] Eric Trist, *The Evolution of Socio-Technical Systems: A Conceptual Framework and an Action Research Program*, Issues in the Quality of Working Life, Occasional Paper No. 2 (Ontario Quality of Working Life Centre, Ontario Ministry of Labour, 1981), 8. Trist's sentence in full reads "Cooperation between task groups was everywhere in evidence; personal commitment was obvious, absenteeism low, accidents infrequent, productivity high." Earlier drafts of this chapter quoted it without the two verbs and with a comma for the semicolon; we have restored Trist's wording. The paragraph immediately before it describes the Haighmoor work organization as "a set of relatively autonomous groups interchanging roles and shifts and regulating their affairs with a minimum of supervision," which is what we paraphrase above.

[9] E. L. Trist, G. W. Higgin, H. Murray, and A. B. Pollock, *Organizational Choice: Capabilities of Groups at the Coal Face under Changing Technologies; The Loss, Re-Discovery and Transformation of a Work Tradition* (Tavistock Publications, 1963).

cannot be designed separately, and it was made from a coal face rather than from a diagram.

For us the consequence is a change of category rather than a change of belief. Read casually, the finding sounds like something the model already says: of course the operational units need enough freedom to handle their own disturbances. Read as its authors meant it, it says something stronger and less comfortable. The metasystem does not hand the amount of freedom down after the plant has been laid out. Choices made when the technology was specified fix it in large part, and the people who made them were not thinking about autonomy at all. Where a seam is worked by a single interchangeable group, the group can absorb its own variance. Where the same seam is worked by a shift-segmented line, no amount of enlightened supervision restores that capacity. Beer saw part of this. "The loss of autonomy, the erosion of liberty, is already implicit in the total system," he wrote, and it goes unrecognized and gets blamed on autocratic corporate management. But he treated the linkages between operations as given, arising "by happenstance" from the nature of the enterprise.[10] At a coal face they are a property of the work system, designed in or designed out before anyone draws the diagram.

We can test that claim and have not. It also warns us about our own practice. A viable system diagram drawn over a work system whose technical design has already destroyed the operational units' capacity to absorb variance will be internally consistent and practically empty.

The method can be separated from the finding, and it matters at least as much. Trist and Bamforth went and looked, at a place where the arrangement had emerged without anyone designing it, and then described it precisely enough that other people could argue with the description. The comparison work came later and was slow. Neither step is beyond us; we have simply not taken them often. Our tradition has case reports written by the consultant who did the intervention, a different genre with different failure modes, and very little observation of arrangements that arose without us.

Herbst and Cherns: the design rule

What they found · Minimum critical specification. **When** · 1974–1987. **Who cited whom** · The rule developed inside sociotechnics, independently of a System Three job description. **Evidence base** · Design principles and case work, without a validated measure of how much specification is critical. **What it costs us** · The missing instrument leaves over- and underspecification as retrospective judgments.

Philip Herbst wrote "Designing With Minimal Critical Specifications" in 1974, inside the same tradition, and the design rule is his: specify no

[10] Beer, *Heart of Enterprise*, 152 and 142; Beer, *Diagnosing the System*, 84, where the operational and environmental linkages "absorb whatever variety the circumstances of the enterprise dictate."

more than is absolutely essential, and do the work of identifying what is essential.[11] Two years later Albert Cherns systematized it as the second of nine principles of sociotechnical design. He was scrupulous about the debt, saying the principles "owe more to the writings of others... than to our own originality."[12] His formulation is the one to carry: "The negative simply states that no more should be specified than is absolutely essential; the positive requires that we identify what is essential."[13]

That is System Three's job description, written three years before *The Heart of Enterprise* and with no knowledge of it. In one respect it also goes further than Beer. The VSM↗ tells System Three what it is *for*: hold cohesion while spending as little of System One's autonomy as possible. Minimum critical specification tells it what to *do*, starting from what is essential.

The two are easy to read as the same thing in different words, but a purpose and a procedure are not interchangeable. Beer did write rules close to Herbst's. The metasystem "*should make minimal use of variety attenuators*" in its dealings downward; you organize a system "only somewhat" and "ride on the dynamics of the system"; the management paperwork is sorted into mandatory red and anti-oscillatory green, so that everyone sees "that a tiny proportion of the total paperwork machinery is red."[14] These rules classify constraints by function. They give no procedure for finding the essential ones. Beer's principle of autonomy is a constraint on outcomes:[15] whatever System Three does, the residue of autonomy left to System One should be as large as cohesion allows. Herbst's is a rule of construction.

[11] P. G. Herbst, *Socio-Technical Design: Strategies in Multidisciplinary Research* (Tavistock Publications, 1974). The chapter "Designing with Minimal Critical Specifications" was reprinted in Eric Trist and Hugh Murray, eds., *The Social Engagement of Social Science: A Tavistock Anthology*, vol. 2, *The Socio-Technical Perspective* (University of Pennsylvania Press, 1993), https://doi.org/10.9783/9781512819052-017. Cherns cites Herbst under both forms at once — "Herbst, 1974/Vol. II" — which is one reason the attribution has been easy to lose.

[12] Chapter 10, Six Components, Six Literatures, No Beer, attributes minimum critical specification to the sociotechnical tradition generally, which is right, and this chapter narrows it: the principle is Herbst's, and Cherns named and systematized it while explicitly disclaiming originality. Cherns's sentence, in full, is that the nine principles "represent a distillation of experience and owe more to the writings of others (Emery, Vol. II, 'The Nine-Step Model'; Emery and Trist, 1972; Herbst, 1974/Vol. II, 'Designing With Minimal Critical Specifications') than to our own originality. They have not, however, previously been systematized." The distinction matters only because Cherns is the citation everyone uses.

[13] Albert Cherns, "The Principles of Sociotechnical Design," *Human Relations* 29, no. 8 (1976): 783–92, https://doi.org/10.1177/001872677602900806. Minimal critical specification is Principle 2 of the nine, and the quoted formulation opens Cherns's discussion of it.

[14] Beer, *Heart of Enterprise*, 157–58 and 180–81; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 53: "Instead of trying to organize it in full detail, you organize it only somewhat; you then ride on the dynamics of the system in the direction you want to go." See also *Heart of Enterprise*, 340 ("no uniform planning rubric can be imposed") and 344.

[15] Beer, *Heart of Enterprise*, 158, for what Beer calls "the cybernetic argument for autonomy (or freedom)", which is what this book means by the principle of autonomy: the metasystem "should make only that degree of intervention that is required to maintain cohesiveness"; 160, where minimal intervention is "essential to viability, save only that this minimum is compatible with systemic cohesion"; 173, where autonomy is the "institutional synonym" of freedom. Beer does not count it among his Principles of Organization; Chapter 16 uses the same definition.

When you write the specification, write the smallest one that closes the essential variances, and treat every additional clause as a cost to be justified. A person drafting a document on a Tuesday can follow the second. The first cannot be followed that way. Practitioners who have tried to apply the principle of autonomy directly will recognize the difficulty; usually the principle is invoked to defend a setting that was arrived at some other way.

Cherns revised the list in 1987, and for a cybernetic reader the revisions are the interesting part. He added *power and authority* and *transitional organization*.[16] Those are precisely the two things on which sociotechnical implementations, and VSM implementations after them, most reliably fail. The pattern is familiar enough to be a diagnosis. A structural design is agreed, the units are given nominal discretion, and the authority to commit resources stays where it was. Or a design is agreed that would work once installed, and no one has specified the organization that carries the work from the present arrangement to that one. Cherns did not have our vocabulary for either failure. He had watched both happen often enough to promote them into principles.

There is a further reason to want the rule alongside the principle. The principle of autonomy is stated in a form that makes it very hard to be caught violating. Any restriction on a unit can be defended as required by cohesion, and any grant of discretion as the residue the principle demands, because the principle states the boundary as an equation that nobody has measured. A rule of construction is less accommodating. Of each clause in a specification it asks which essential variance the clause exists to close, and it treats a clause that cannot answer as surplus. A design review can actually put that question, and losing it is part of what the model's silence costs.

The rule has a limit: it remains a rule and not a measurement. Neither Herbst nor Cherns produced a way of telling, in advance, how much specification is critical. Their tradition left that to skilled joint design with the people doing the work, which is a defensible answer and not an instrument. We have no instrument either, and we have less excuse, because the question is stated in exactly the terms our tools were built for.

What the sociotechnical record will and will not support

The tradition also has evidence, and more of it than we do. In 1982 Pasmore, Francis, Haldeman and Shani reviewed a hundred and thirty-four sociotechnical experiments and reported them successful across productivity, cost, quality and satisfaction. The same review indicts itself: very few

[16] Albert Cherns, "Principles of Sociotechnical Design Revisited," *Human Relations* 40, no. 3 (1987): 153–61, https://doi.org/10.1177/001872678704000303. The revised list keeps compatibility, minimal critical specification, information flow, support congruence and incompletion, folds variance control and the multifunctional principle back in, drops *design and human values* as a separate item on the ground that values underwrite all the principles rather than standing beside them, and adds *power and authority* and *transitional organization*.

of the hundred and thirty-four involved redesigning the *technology*. Most rearranged the social system around fixed plant, which is not what the theory claims to do.[17]

That qualification matters more than the headline, and our own literature rarely makes one like it about itself. A hundred and thirty-four cases that mostly tested half the theory do not validate joint optimization. They validate social rearrangement under fixed technical constraint, a narrower and more familiar claim. The authors said so in print, in the same paper that reported the successes.

John Kelly's 1978 reappraisal went further. Autonomy in self-managing groups, he argued, is bounded by economic imperatives rather than genuinely devolved; joint optimization is a myth better described as labor intensification; and pay incentives carry more of the causal weight than the theory admits.[18] Kelly's argument should be taken as a live hypothesis rather than a historical footnote. It is a rival explanation of the same results, and it applies to us without modification. If devolved discretion in these cases was bounded by economic imperatives set elsewhere, the experiments demonstrated the productivity of a particular reorganization of effort, and not the value of autonomy. Any structural design that claims to devolve discretion while leaving the money and the headcount where they were is exposed to the same reading, and we have not built the study that would distinguish the two.

And then the tradition died without being refuted, which is the part cybernetics should study hardest. The oil shock, the political turn of the eighties, management's preference for control-oriented alternatives and union indifference to shop-floor initiatives all played a part. The decisive cause was an action-research method that could never diffuse beyond its pilot sites. Harry Scarbrough's verdict, quoted approvingly in a recent reassessment of the tradition, is that it became "a kind of intellectual Esperanto: a language championed by a brave and enthusiastic few, but actually spoken by no one else."[19]

---

[17]William Pasmore, Carole Francis, Jeffrey Haldeman, and Abraham Shani, "Sociotechnical Systems: A North American Reflection on Empirical Studies of the Seventies," *Human Relations* 35, no. 12 (1982): 1179–1204, https://doi.org/10.1177/001872678203501207. The self-indictment is the authors' own: sociotechnical experiments have been "extremely successful overall," but "the number of experiments involving technological innovation or change is relatively small," and they concede that "we still have much to learn regarding the design of technical systems for joint optimization."

[18]John E. Kelly, "A Reappraisal of Sociotechnical Systems Theory," *Human Relations* 31, no. 12 (1978): 1069–99, https://doi.org/10.1177/001872677803101204. Kelly argues that the autonomy granted to autonomous work groups is "limited, and subordinate to particular economic objectives," that joint optimization "has little connection with the reality of sociotechnical practice" and is better described as labor intensification, and that "the role of pay incentives is also shown to have been seriously underestimated."

[19]H. Scarbrough, "The Social Engagement of Social Science: A Tavistock Anthology," *Human Relations* 48, no. 1 (1995): 23–33, https://doi.org/10.1177/001872679504800102. We have not been able to confirm the "intellectual Esperanto" sentence against the text of Scarbrough's review-article, which is paywalled, nor to identify with certainty the recent reassessment that

That sentence describes a risk we share. The sociotechnical tradition had field results, a design rule, a review of a hundred and thirty-four cases and a serious internal critic, and it still did not diffuse. On each of those counts our position is weaker. Whatever protected the tradition from being wrong did not protect it from being unused, and nothing in our situation suggests we are better placed.

Endenburg: the mechanism

What he found · Consent, circles and the double link as a structural response to autonomy–cohesion. **When** · Boeke's consent principle in 1945; Endenburg's later engineering synthesis. **Who cited whom** · Romme's 1995 scholarly treatment has eight references and no cybernetician. **Evidence base** · Operating organizations, but no reliable adoption census. **What it costs us** · The double link answers a mechanism question Beer posed but did not settle.

In the nineteen-seventies, the decade of *Brain of the Firm*, Gerard Endenburg built the sociocratic circle method inside his family's electrical engineering firm. It has four rules. Decisions stand when no member raises an argued objection. Work is done in semi-autonomous circles, each with its own aim, deciding by consent. Adjacent circles are linked by *two* people, a leader appointed downward and a delegate elected upward, both full members of both circles. And people are assigned to functions by consent after open argument.[20]

The double link is the sharpest mechanism in this chapter. Beer's recursion↗ requires that command and resource channels between levels have adequate variety and does not say how to build them. Endenburg implements the requirement as a membership rule and makes the upward channel *elected*. That is a variety design and a legitimacy design at once. Beer multiplied the upward channels: six vertical components, an audit that goes directly to the operations, an algedonic bypass. He put the heads of System One, and perhaps "simply workers," into System Five.[21] He never made the upward link an elected membership with a right of objection. The double link adds election and consent to a channel structure Beer already required to be plural.

In our own terms the design repays a closer look. A single link between levels, which is what most organizations have and what our diagrams imply

quotes it; the sentence circulates in the sociotechnical literature without a page reference. The review-article itself exists and is correctly described. Anyone reprinting the quotation should check it against page proof before doing so.

[20] Gerard Endenburg, *Sociocracy: The Organization of Decision-Making; "No Objection" as the Principle of Sociocracy* (Eburon, 1998). For a compact statement of the four rules see also A. Georges L. Romme and Gerard Endenburg, "Construction Principles and Design Rules in the Case of Circular Design," *Organization Science* 17, no. 2 (2006): 287–97, https://doi.org/10.1287/orsc.1050.0169.

[21] Beer, *Heart of Enterprise*, 143, 263–64 and 407; Beer, *Diagnosing the System*, 82; Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 149.

when they are read carelessly, gives the higher circle exactly one account of what is happening below, produced by the person whose performance is being reported. Two links do more than double the bandwidth. The second occupant is selected by the people being described and sits in the higher circle as a full member with the power to withhold consent, so the channel structurally cannot carry only one interested account. Whether that raises the effective variety of the upward channel or mainly its credibility is precisely the question nobody has tested, and the two are different things.

As far as we can establish, the method is also not cybernetic in citation. Every account of its origin traces it to Kees Boeke's consent principle, itself out of the Quaker business method, and to feedback and control as Endenburg had learned them as an engineer.[22] The earliest scholarly treatment, Georges Romme's in 1995, names no cybernetician at all: not Ashby, not Wiener, not von Foerster, not Beer.[23] This is the same finding as the Galbraith case in Chapter 12, *The Variety Logic*, and the four silences in Chapter 11, *The Nesting: Simon, Koestler, Jaques*, in a fourth place.

Consent is easy to hear as a cultural preference and deserves the same structural reading. The rule that a decision stands unless a member raises an argued objection is a filter on what has to reach the center at all. It admits the objection and discards the preference, and it does so at the level where the objection arises, before aggregation. That makes it a variety-attenuating device in the strict sense, placed where attenuation is cheapest. Whether it attenuates the right things, and what it discards along with the noise, is an empirical matter on which the sociocratic literature is thinner than its confidence would suggest.

Four independent arrivals do not prove that the components are correct. They are evidence that the components are findable, a weaker claim and a more useful one. They also mean that the case for our model cannot rest on having the components, since other people have them. Whatever we are entitled to claim has to rest on the assembly, and Part Four takes up that argument.

One hard institutional fact

The Dutch Social and Economic Council, the country's tripartite advisory body, will exempt a company from the statutory obligation to establish a works council, and full implementation of the sociocratic circle method is

[22] Kees Boeke, *Sociocracy: Democracy as It Might Be* (1945), https://www.sociocracy.info/sociocracy-democracy-kees-boeke/.

[23] A. Georges L. Romme, "The Sociocratic Model of Organizing," *Strategic Change* 4, no. 4 (1995): 209–15, https://doi.org/10.1002/jsc.4240040404. The claim is about citation, not vocabulary: Romme's article does use the word *cybernetics* and does describe circle processes as feedback loops, but it sources that language to Endenburg's own engineering training by way of John Buck and Gerard Endenburg, *The Creative Forces of Self-Organization* (Sociocratisch Centrum, 1984). Its reference list contains eight items and not one cybernetician among them.

one of the accepted grounds.[24] It is the strongest hard fact in the sociocratic literature, and it is legal rather than empirical. As far as we have been able to find, no model derived from cybernetics has been accepted anywhere as a substitute for statutory co-determination. The same fact also caps the claim: the companies currently holding that exemption on sociocratic grounds number in single figures.[25]

The exemption matters for a specific reason. It is a judgment, by a body with an interest in employee representation, that a particular structural design delivers what the statute was written to deliver. That is an external assessment against a stated standard, and no VSM implementation has ever had to pass one. It is not evidence that the method performs well on any other dimension, and the small numbers mean it cannot bear weight as evidence of diffusion. What it establishes is that a governance mechanism of this family can be made legible enough, and specified tightly enough, for a regulator to rule on it. Ours has never been specified to that standard.

The descendants

The descendants are where the evidence gets uncomfortable. Holacracy, formalized by Brian Robertson in 2007, took the circle structure and the name from Koestler's holarchy.[26] Zappos adopted it across a workforce of roughly fifteen hundred. Fourteen percent of staff took a 2015 buyout offer, departures reached roughly eighteen percent within a year, and the company moved off the system over the following four.[27] The reported reason belongs in a cybernetic book. Corporate Rebels reports Tony Hsieh

[24]Sociaal-Economische Raad, "Ontheffing van de plicht een or in te stellen," https://www.ser.nl/nl/thema/medezeggenschap/medezeggenschapstructuren/ontheffing. The power is under article 5 of the Wet op de ondernemingsraden. The SER's own wording is that a good ground for exemption may be that the employees are substantially co-owners "of dat de Sociocratische Kringorganisatiemethode volledig is geïmplementeerd" — that the sociocratic circle organization method has been fully implemented. Two qualifications the text does not carry: an exemption is temporary, granted for at most five years and renewable on reapplication, and the SER states that it is sparing with them.

[25]SER, page cited in the previous note, section "Overzicht verleende ontheffingen" (consulted 29 September 2026). At the time of writing the SER's published register lists six current exemptions in total, of which four are granted on sociocratic grounds: Talent B.V. (2023), Endenburg Elektrotechniek B.V. (2023), CKO KleurRijk (2024) and Schuberg Philis (2024). The other two — BügelHajema Adviseurs B.V. (2021) and Holland Colours N.V. (2023) — rest on employee self-governance and a works-council-substitute committee rather than on the circle method. "Single figures" is therefore if anything generous.

[26]Brian Robertson, "Organization at the Leading Edge: Introducing Holacracy; Evolving Organization," *Integral Leadership Review*, June 2007, https://integralleadershipreview.com/5328-feature-article-organization-at-the-leading-edge-introducing-holacracy-evolving-organization/; Brian J. Robertson, *Holacracy: The New Management System for a Rapidly Changing World* (Henry Holt, 2015); HolacracyOne, *Holacracy Constitution, Version 5.0* (HolacracyOne, 2019), https://www.holacracy.org/constitution. Robertson credits the term to Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), where the holon is introduced as a whole that is also part of a larger whole.

[27]For the headcount and the 2015 offer, Benjamin Snyder, "14% of Zappos' Staff Left after Being Offered Exit Pay," *Fortune*, May 8, 2015, https://fortune.com/2015/05/08/zappos-quit-employees/ — roughly 210 of about 1,500 staff, or 14 per cent, took at least three months'

as saying that the circles were still arranged hierarchically, with budget and headcount allocated from the top down, so teams turned inward instead of toward customers.[28]

In Beer's exact terms, that is a System Three failure. The governance layer was replaced and the resource channel was not, and the resource channel won. How much this shows needs care. The wording is second-hand, as the note records, and one executive's retrospective account of a reversion is not a diagnosis. What can be said with more confidence is that the observable design put a decentralized decision procedure on top of a centralized allocation of budget and headcount, and that the model predicts a specific consequence for that combination. Units whose discretion does not extend to resources will spend their discretion on the things it does extend to, which are internal. The prediction and the reported outcome agree. Nobody has done the work that would turn an agreement into a finding.

Medium abandoned Holacracy too, and its stated second reason is the sharper one: "the act of codifying responsibilities in explicit detail hindered a proactive attitude."[29] At Medium, Holacracy failed by breaching Cherns's second principle, a rule written more than thirty years before Holacracy existed by the tradition Holacracy's own family tree descends from. Neither community appears to have connected the two.

Taken together, the two cases make a small argument for the design rule rather than against the family of designs. Zappos left the specification too thin in one place, at the resource channel; Medium made it too thick everywhere. Both failures can be described in the vocabulary of minimum

---

severance by the April 30, 2015 deadline. For the year's total, Jena McGregor, "Zappos Says 18 Percent of the Company Has Left Following Its Radical 'No Bosses' Approach," *Washington Post*, January 14, 2016 — 260 people, 18 per cent, over 2015. See also Bebeto Samantarai and Indu Perepu, *Holacracy at Zappos (B): From Experiment to Exodus* (Case Centre, 2017), https://www.thecasecentre.org/products/view?id=146418. The "following four" years is the interval over which Zappos left the system rather than a single announced exit: the internal marketplace that displaced the circle structure was introduced in March 2017, customer-generated budgeting in September 2018 and the triangle of accountability in early 2019, and the company had quietly stopped describing itself as holacratic by 2020. The three dates are those given in Minnaar, "Zappos's Evolution" (next note). Zappos has never published a statement abandoning Holacracy, and we have found none.

[28] Joost Minnaar, "Zappos's Evolution: From Holacracy to Market-Based Dynamics," Corporate Rebels, October 2, 2021, https://www.corporate-rebels.com/blog/zappos-market-dynamics. The sentence as printed there is "Our circles were still arranged hierarchically, where budgeting or headcount was allocated from the top down." Two cautions. Corporate Rebels gives no interview, transcript or date for the remark, and we could not find it in any Zappos publication, including the Zappos Insights page on Holacracy to which the quotation is elsewhere attributed; Hsieh died in November 2020 and cannot be asked. The gloss that follows — that this turned teams inward rather than toward customers — is Corporate Rebels' summary of the consequence and not Hsieh's own words.

[29] Andy Doyle, "Management and Organization at Medium," Medium, March 4, 2016, https://medium.com/blog/management-and-organization-at-medium-2228cc9d93e9. Doyle's sentence in full: "we found that the act of codifying responsibilities in explicit detail hindered a proactive attitude and sense of communal ownership." His first reason is coordination cost at scale; the sentence we quote is the second.

critical specification, and both in ours. Neither community had an instrument that would have shown the error while there was still time to correct it. That missing instrument, more than either company's judgment, is the common cause worth pursuing.

What formal equality does not reach

The best critical study we found looked at a co-housing community through sixteen interviews. It reports real decentralization and reduced coercion alongside four inequalities that formal equality does not touch: founders keep status, members hold divergent expectations of commitment, circles apply the rules inconsistently, and people differ in time, procedural knowledge and emotional skill.[30] Beyond that, the peer-reviewed evidence for sociocracy itself is thin, and its own advocates say so.

That is the precise limit. Sociocracy equalizes the decision *procedure* and leaves the distribution of *capacity* exactly where it was. That is a variety problem, stated in a vocabulary that is not ours, and nobody has modeled it. One site and sixteen interviews cannot establish how general the finding is, and a residential community is not a firm. What the study does supply is a well-specified hypothesis: a consent rule redistributes the right to object without redistributing the ability to formulate an objection, and the second distribution determines who is actually heard. Anyone who has watched a well-run participative process converge on the preferences of its most articulate members will recognize the claim. Testing it is within reach of our tools and has not been attempted.

What the numbers actually say

Earlier drafts of Chapter 10, *Six Components*, claimed that Endenburg's system is installed in more organizations today than the VSM is. No figures support that claim, and the honest position is narrower.

Nobody counts either. Endenburg's own institute claims hundreds of adopters across the nineteen-eighties and nineties and offers no current total.[31] One sociocratic body reports direct knowledge of about a hundred

[30] Daniel King and Martyn Griffin, "Governing for the Common Good: The Possibilities of Sociocracy in Nonprofit Organizations," *Voluntas: International Journal of Voluntary and Nonprofit Organizations* 35, no. 6 (2024): 1046–57, https://doi.org/10.1007/s11266-023-00627-2. The site is a co-housing community of twenty-five households in the northwestern United States, studied through sixteen interviews and four days of observation including four general circle meetings. The four inequalities are the authors' own categories: status, expectation, application and capacity.

[31] The Sociocracy Group, successor to Endenburg's Sociocratisch Centrum, states on its website that over the 1980s and 1990s "hundreds of organizations in the Netherlands and around the world adopted Sociocracy" (The Sociocracy Group, "About Us," https://thesociocracygroup.com/about-us/, checked 29 September 2026). It gives no list, no count and no current total, and neither does Sociocratisch Centrum Nederland. The earliest scholarly description, Romme, "The Sociocratic Model of Organizing," 211, names only a handful of adopters by kind — Endenburg Elektrotechniek, a municipal police department, a consultancy, a chain of

organizations and assumes there are three times more, an estimate wearing the clothes of a count.[32] Holacracy publishes an opt-in list and, beside it, a claim of thousands of users with no total.[33] Sociocracy 3.0 is openly licensed and by design keeps no registry at all.[34] On our side there is no registry either. The most recent systematic review of the VSM screened seven hundred and sixty-eight documents, found twenty-one studies meeting its inclusion criteria, and concluded that most of what exists is model and simulation rather than implementation, with limited empirical validation and almost no longitudinal work.[35]

The comparison is between two unaudited numbers, one plausibly in the hundreds and the other, in the documented literature, in the tens. It points one way and it is not a census.

Two things follow. The weaker is that the claim should not be made without a count, and Chapter 10 no longer makes it. The stronger is that the absence is symmetrical and therefore diagnostic. Neither community can say how many organizations are running its design, and neither has treated that as a defect worth fixing. For a tradition whose central claim is that organizations can be designed for viability, not knowing how many designed organizations exist is a peculiar thing to have tolerated for fifty years.

As of September 2026

The numbers in this section move, and one of them moves fast. The Dutch Social and Economic Council's published register of works-council exemptions lists six in total, of which four rest on sociocratic grounds: Talent B.V. and Endenburg Elektrotechniek in 2023, CKO KleurRijk and

hairdressing shops — and adds that the model "has been tested in organizations of up to 1300 people and is being tried out on a limited basis in substantially larger organizations in the Netherlands, Germany, Brazil, Canada, the USA and other countries." That is compatible with hundreds and does not establish it.

[32] Sociocracy For All, "Sociocracy: The Movement," https://www.sociocracyforall.org/sociocracy-the-movement/. The organization states that it "is aware of and connected to about 100 sociocratic organizations" and assumes "there are at least 3 times more organizations that have implemented sociocracy to some significant degree," while conceding that "it is impossible to make a realistic guess of how many organizations are using sociocracy worldwide."

[33] HolacracyOne, "Directory," https://www.holacracy.org/directory/, checked 29 September 2026: "This directory lists organizations that choose to be featured," and "Thousands of organizations around the world use Holacracy, but only the ones who've opted in are listed here."

[34] Bernhard Bockelbrink, James Priest, and Liliana David, *Sociocracy 3.0 — A Practical Guide* (Evolving Collaboration, 2022), https://patterns.sociocracy30.org/. The pattern library is published under a Creative Commons license; there is no adoption registry by design.

[35] Mohammad Reza Fathi, Soraya Birami, Alireza Payvar, and SeyedAli Doorafshan, "A Systematic Review of the Viable System Model: Applications, Insights, and Future Directions," *Journal of Systems Thinking in Practice* 4, no. 3 (2025): 109–45, https://doi.org/10.22067/jstinp.2025.91552.1136. The seven hundred and sixty-eight is the Scopus record count; a further forty-five records came from manual searching, for eight hundred and thirteen after deduplication, and twenty-one studies met the inclusion criteria. The authors report that "most studies rely on simulations or conceptual models without real-world testing" and that there is an "absence of longitudinal studies."

Schuberg Philis in 2024. The systematic review of the Viable System Model used for the comparison screened seven hundred and sixty-eight Scopus records, eight hundred and thirteen after manual searching and deduplication, and admitted twenty-one studies.

None of these is an adoption count, and neither side of the comparison has one. The register this chapter asks this community to keep does not exist on the other side either. Current values, and the dates they were taken, belong in the companion register rather than in a printed chapter — which is the argument of the chapter that follows this one.

What this chapter hands forward

Operationalize minimum critical specification. Fifty years on there is still no measure of how much specification is critical, and no way to tell over- from under-specification before the fact. That is a variety-engineering question. Cybernetics owns the tools for it and has never built the instrument. Medium over-specified and Zappos left the resource channel where it was; the two failures are different, and anyone who bothers can diagnose both in Beer's vocabulary. The instrument need not be elegant to be useful. A defensible first version would be a procedure for counting the essential variances a specification is supposed to close, and for identifying the clauses that close none of them.

Test the double link. Does an elected upward delegate actually raise the metasystem↗'s variety, or mainly confer legitimacy? Either answer would be informative. If the gain is in variety, the double link belongs in our recommended designs as a channel specification. If it is in legitimacy, it belongs there anyway, for a different reason, and our account of what upward channels are for needs revision. The mechanism has been running in real organizations for fifty years without anyone asking.

Write up Zappos. It is the cleanest natural experiment this tradition has been handed since Chile: governance decentralized, resource allocation left centralized, and the system reverted for a reason its own chief executive stated in System Three's terms without knowing the name. Nobody has done it. The documentary record is public, the sequence of changes is dated, and the rival explanations can be stated in advance.

The counter-entry: sociotechnics did not win

Kelly's 1978 criticism that autonomy remained subordinate to economic objectives, and Scarbrough's 1995 obituary, belong in the ledger. They cost the chapter any easy claim that an independently discovered design rule naturally diffuses. A rule can be correct, published, systematized, taught and reviewed across a hundred and thirty-four cases, and still leave almost no trace on how work is organized. If the argument of this part of the book were that good components win, the sociotechnical record would refute it. The argument is narrower: these components were found without us,

which tells us something about the components and nothing at all about their prospects.

The priced receipt

Write the Zappos natural experiment from the existing documentary record: a bounded desk study of the 2017–2020 reversion, governance changes, and the retained resource-allocation channel. Budget ten researcher-days and one legal/source check. Pre-state the rival explanations. If central resource allocation does not explain the reversion better than the alternatives, surrender the claim that the failure was a clean System Three diagnosis. The cheaper study comes first because the documents already exist.

Chapter 14

The Evidence: Ostrom

What she found · Eight design principles for durable common-pool-resource institutions. **When** · Coding manual in 1989, synthesis in 1990, later diagnostic turn through 2010. **Who cited whom** · This is an independent route to multi-level governance. **Evidence base** · More than a thousand collected cases, a 455-question instrument, 86 coded cases, and three later reanalyses. **What it costs us** · Nesting is the least-supported principle in the 2010 review, and this literature published disconfirmations the VSM literature has not equipped itself to produce.

An independent route, not a rival program

Irrigation systems, inshore fisheries, mountain commons, groundwater basins: Elinor Ostrom wanted to know why some of them survive for six hundred years and others collapse in a generation. She was not studying firms. *Governing the Commons* appeared in 1990; the Nobel followed in 2009.[1]

We have cited her for a long time, and warmly. Nested, multi-level governance, arrived at from irrigation ditches by a scholar whose cybernetics came from Ashby rather than from Beer,[2] is among the better reasons to think the model's components are not arbitrary. Chapter 10, *Six Components*, makes that argument, and this chapter does not withdraw it. What we have not done is read the rest of her literature with the attention we gave the part that flattered us. Her tradition wrote down what it was looking for before it looked, coded a large number of cases against that written instrument, submitted the result to independent reanalysis three times, and printed the findings that went against it. Two of those findings go against us as well. We cited this literature as corroboration and did not report where it goes against us. This chapter repairs that omission.

[1] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990), https://doi.org/10.1017/CBO9780511807763. The prize is the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel, shared with Oliver Williamson in 2009.

[2] Paul Lewis and Paul Dragos Aligica, "The Ostroms on Self-Governance: The Importance of Cybernetics," *Journal of Institutional Economics* 20 (2024): e23, https://doi.org/10.1017/S1744137424000079. The article traces the Ostroms' multiple levels of rules to Ashby's argument that self-organization needs a hierarchy of feedback mechanisms, largely by way of Vincent Ostrom, who had talked at length with Ashby at Stanford in 1955–56. It does not mention Beer.

Ostrom's program is no opponent, and the chapter keeps no score. Her literature is the nearest available example of what a comparable body of work looks like when a community decides to gather evidence about its own claims. The comparison with our own record is unflattering to us, and the unflattering parts should be stated precisely, neither dramatized nor waved through.

The eight principles, read with the model in hand

Her eight design principles for long-enduring common-pool resource institutions read as follows. Clearly defined boundaries. Congruence between the appropriation and provision rules and local conditions. Collective-choice arrangements, in which most of those affected by the operational rules can participate in changing them. Monitoring by monitors who "are accountable to the appropriators or are the appropriators". Graduated sanctions. Rapid, low-cost, local conflict resolution. Minimal recognition by external authorities of the right to organize. And, for resources embedded in larger systems, nested enterprises, with appropriation, provision, monitoring, enforcement, conflict resolution and governance "organized in multiple layers".[3]

Four of the eight land on our own territory, and we put none of them there. The eighth is nesting, reached from irrigation ditches rather than from Beer's derivation. It resembles recursion and is not the same claim. Beer's recursion requires each embedded layer to be a viable system in its own right; Ostrom's principle asks that activities be "organized in multiple layers of nested enterprises."[4] The eighth principle is convergence on a related but not identical object. The third is a legitimacy condition on System Five↗: those who live under the operational rules can take part in changing them. That constrains where policy may be made. Beer comes near it: System Five "consists of many people," should include workers, and later represents all stakeholders, and divisions should preferably take

[3]Ostrom, *Governing the Commons*, 90 (table 3.1). Principle 4 reads in full: "Monitors, who actively audit CPR conditions and appropriator behavior, are accountable to the appropriators or are the appropriators." Principle 8 reads: "Appropriation, provision, monitoring, enforcement, conflict resolution, and governance activities are organized in multiple layers of nested enterprises," and Ostrom states it only "for CPRs that are parts of larger systems."

[4]Beer's criterion is set out in Chapter 10: Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2, where recursion is not "any kind of system contained inside another"; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 308, where it must not become "a meaningless appendage to the term *organizational level*." Ostrom's wording is quoted in the note to her eight principles above.

part in formulating objectives.[5] But he files participation under "the psychology of the system (and not its logic)"; Ostrom makes it a condition of endurance. The fourth is the audit channel (the Three-star of *Diagnosing the System*) with an ownership requirement Beer did not impose, since a monitor accountable to the appropriators is a different instrument from a monitor accountable to the center. Beer's audit runs "by agreement with System One management,"[6] which protects the audited unit's consent and does not make the auditor accountable to it. Chapter 4, *Channels Without Owners*, takes that requirement up, because our specification of the audit channel says who may open it, how often and with whose agreement, and says nothing about whose interest the auditor serves. The seventh is a precondition the VSM places and does not regulate. In the model, the right to self-organize lies with the containing recursion, whose metasystem has "the authority to delimit the autonomy of any … System One" by the cybernetic rules.[7] The model has no case in which that authority dissolves the system, or is held by a power the cybernetic rules do not bind. Chapter 8, *The Loop That Stayed an Intention*, uses it as a boundary condition on the whole apparatus. Chile in September 1973 is the proof of Ostrom's point about Beer.

The other four do not map, and that is informative too. Graduated sanctions and rapid local conflict resolution are not variety-engineering constructs, and the model offers no guidance on how to build either. Clearly defined boundaries and congruence with local conditions come closer to the recursion criteria we do discuss, but they are stated as properties of a rule set rather than of a system's structure. Independent arrival should produce a partial overlap. A complete one would more likely mean that we had read the mapping into her table than found it there.

How the evidence was built

The Workshop's coding manual was built by evaluating questions against more than a thousand unpublished case studies, and it runs to four hundred and fifty-five coding questions. Eighty-six cases were systematically coded. Fourteen are read in depth in the book, and five of those fourteen are failures.[8] That evidence is what makes this material belong in a book about our own poverty. Including the failures makes the principles

[5]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 263–64; Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 149; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 159, where objectives "must, at the very least, be explained carefully to the divisions; preferably the divisions should participate in their formulation," and participation belongs to "the psychology of the system (and not its logic)."

[6]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 82.

[7]Beer, *Brain of the Firm*, 275; Beer, *Diagnosing the System*, 1, on the viable system that "cannot survive in a vacuum"; Beer, *Heart of Enterprise*, 264.

[8]Elinor Ostrom, Arun Agrawal, William Blomquist, Edella Schlager, and S. Y. Tang, *Common-Pool Resources (CPR) Coding Manual* (Workshop in Political Theory and Policy Analysis, Indiana University, 1989), https://seslibrary.asu.edu/resources/. The figures of 455 coding

falsifiable rather than an anthology of successes. Beer argued for the idea: he proposed "organizational pathologists" to do post-mortems on systems that are viable no more, and he reported failures of his own.[9] The VSM literature has never included failures in a coded sample.

The instrument matters as much as the count. Because the coding questions were written down and published, a later researcher could recode the same cases, disagree with the original judgments in public, and show her work. That is how the reanalyses below became possible, and it is why the numbers in this chapter can be argued with at all. Our own case literature has no equivalent artifact. Beer published a diagnostic workbook whose charts readers may reproduce, and a fifteen-question identity instrument with card scoring and clustering.[10] Neither was used as a coding instrument across cases, neither carries success criteria stated in advance, and nobody has recoded a case with either. So a skeptical reader has nothing to recode, and nobody outside a given consultancy can establish that a reported success was a success by a standard fixed before the engagement began.

How the evidence was tested

Then the principles were tested. This is the part that matters most, and the part we have no counterpart to. A 2010 review of ninety-one studies found them moderately supported overall, with a revealing gradient: the empirical studies were more supportive than the abstract ones, and the purely theoretical studies were actively unsupportive.[11] The gradient admits more than one reading. People who watch an irrigation system for a year may see structures that a modeler working from published summaries cannot. Or fieldworkers who know the principles may simply find them. The review does not settle which, and neither do we.

---

questions and more than one thousand unpublished case studies behind their construction are not stated in the manual itself; they come from Michael J. Bernstein et al., "Mapping Ostrom's Common-Pool Resource Systems Coding Handbook to the Coupled Infrastructure Systems Framework to Enable Comparative Research," *International Journal of the Commons* 13, no. 1 (2019): 528–52, https://doi.org/10.18352/ijc.904. The manual's own account of its evidentiary base is a bibliography of 4,977 items, over 1,800 archived documents and 1,600 screened and abstracted. We flag one discrepancy we could not resolve: the manual reports fifty irrigation, thirty-two fishery and three other cases coded and entered, which is eighty-five, not the eighty-six usually cited; the difference is probably a later addition to the database, but we have not confirmed that.

[9] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150; Beer, *Brain of the Firm*, 271, 350, 362.

[10] Beer, *Diagnosing the System*, 137; Beer, *Heart of Enterprise*, 417–33.

[11] Michael Cox, Gwen Arnold, and Sergio Villamayor-Tomás, "A Review of Design Principles for Community-Based Natural Resource Management," *Ecology and Society* 15, no. 4 (2010): art. 38, https://doi.org/10.5751/ES-03704-150438. Ninety-one studies and seventy-seven cases were coded. The overall mean evaluation score was 3.73 on a five-point scale, just below the value of four the authors label moderately supportive. The gradient is between empirical and abstract work rather than a clean ranking by fieldwork: detailed case studies scored 3.85, statistical studies 3.50, synthesis studies 4.04, and the nine abstract or theoretical studies 2.33.

A 2016 reanalysis of sixty-nine recoded cases, using qualitative comparative analysis, concluded that "no DP, alone, leads to successful outcomes… but a combination of more than eight DPs seems to greatly increase the odds of success". The principles work configurally, and which configuration works depends on the infrastructure.[12] That substantially qualifies the design-principle idea as it is usually quoted, and it was published by researchers working inside the tradition rather than against it.

A 2017 review of fifty-four papers across every inhabited continent found the two enforcement principles, monitoring and graduated sanctions, consistently weak in practice. It reported that the principles "do not conclusively diagnose the functionality of young and viable CPR institutions, whereas they may do so for either the short-lived (failed) or the long-lasting institutions."[13] Each of the three studies makes the underlying claim smaller and more conditional than the version that circulates in secondary citation. None of the three was suppressed, and the people whose work they qualify cite all three routinely.

Three findings from that literature should make uncomfortable reading for us.

The first finding: nesting is the weakest of the eleven

In the 2010 review the *least* supported of the principles is the nested one, the principle closest to recursion. Nested enterprises came out alone at the bottom of the eleven. That is still within the range the reviewers call moderately supportive, but nesting is the only one of the eleven whose association with success failed to reach significance at the five percent level. The support ratio is 2.0, at the very edge of the weakest band the authors use; Fisher's exact test gives P = 0.065; the effect size, phi = 0.262, is the smallest of the set. Proportionality between what a member puts into a commons and what they take out came out near the top, and it is also one

[12] Jacopo A. Baggio, Allain J. Barnett, Irene Perez-Ibarra, Ute Brady, Elicia Ratajczyk, Nathan Rollins, Cathy Rubiños, et al., "Explaining Success and Failure in the Commons: The Configural Nature of Ostrom's Institutional Design Principles," *International Journal of the Commons* 10, no. 2 (2016): 417–39, https://doi.org/10.18352/ijc.634. Sixty-nine of the Cox, Arnold and Villamayor-Tomás cases were recoded — twenty-five forestry, twenty-four irrigation, twenty fishery — of which twenty-seven were complete on every variable. The two conditions the authors find "key for CPR success independent from the natural and human hard made infrastructure" are 2A, congruence between local conditions and rules, and 2B, proportionality between investment and extraction.

[13] Sirak Robele Gari, Alice Newton, John D. Icely, and Maria del Mar Delgado-Serrano, "An Analysis of the Global Applicability of Ostrom's Design Principles to Diagnose the Functionality of Common-Pool Resource Institutions," *Sustainability* 9, no. 7 (2017): art. 1287, https://doi.org/10.3390/su9071287. Fifty-four papers. The authors report that the principles most commonly present were 1, 2 and 7, that "both the monitoring activities (DP4) and the sanctions for rule violations (DP5) were weak," and — because most of the resources studied were small in scale — that principle 8 was largely not applicable.

of the two conditions the 2016 configural analysis found holding across all three sectors. Nesting held nowhere in particular.[14]

We had counted this place as support. Chapter 10 lists Ostrom's nested governance, arrived at without Beer, among the six independent arrivals that make the model non-arbitrary, and marks it as convergence on a related object. It is an independent arrival. It is also the one her own literature supports least, and both things are true at once. Chapter 5, *The Ladder Without a Ruler*, works with this result at length. The result does not disconfirm recursion. Principle 8 corroborates recursion, and the 2010 result counts against it, only to the extent that Ostrom's nested enterprises are themselves viable systems, and nobody has coded them for that; the second arm of the pilot priced at the end of this chapter would. The further problem it raises for us is that we have no scale on which recursion could be graded in the first place.

Two qualifications are owed here, and they run in our favor without rescuing us. First, a support ratio of 2.0 is a weak positive, not a negative: the cases where nesting was present and the institution endured still outnumber, two to one, the cases where it was present and the institution failed. Second, the 2017 review found principle 8 largely inapplicable to its sample, because most of the resources studied were small in scale. A principle stated only for systems that are parts of larger systems will always be tested on an unfavorable sample when the sample is dominated by systems that are not. What we cannot do is quote the eleven-principle result as evidence for recursion while quietly omitting where nesting placed in it.

The second finding: viability confirmed only in retrospect

The second is the survivorship structure of the whole exercise. The principles were derived from institutions selected for having endured, so endurance is built into the sample before any coding begins. A set of principles that can diagnose the long-lasting and the already-dead, but not the young and viable, describes what survival looks like in retrospect. It offers no instrument for the present tense.

VSM practice makes exactly the same move, and makes it without the coding. Beer did not. "The fact that it is there does not prove that it is

[14]We give the per-principle support ratios reported by the 2010 review from table 3 of the paper itself, as ratios of supportive to unsupportive cases: 4B, accountability of monitors, 11.7; 2B, proportionality, 10.0; 1A, user boundaries, 5.3; 2A, congruence with local conditions, 5.0; 4A, presence of monitors, 3.8; 6, conflict resolution, 3.2; 3, collective choice, 2.8; 5, graduated sanctions, 2.3; 1B, resource boundaries, and 7, right to organize, 2.2 each; and 8, nested enterprises, 2.0. Nesting is therefore alone at the bottom and not, as we had it in an earlier draft, tied with the right to organize. It is also the only one of the eleven for which Fisher's exact test failed at the five per cent level ($P = 0.065$; the authors report it as significant at the 10 per cent level), and it carries the smallest effect size (phi = 0.262). The authors' own bands treat ratios of 2 to under 5 as moderately well supported, so nesting sits at the very edge of the weakest band they use.

effectively there, witness universities, nor efficiently there, witness hospitals," he wrote, and he proposed forward-looking measures: "nobody is measuring stability. It can however be done."[15] The community did not build or test them. In practice we define success as viability, and viability is only ever confirmed afterwards. A firm that is still trading is evidence for whatever diagnosis was applied to it; a firm that failed is evidence that the diagnosis was not implemented properly. Chapter 9, *Why None of This Ever Cost Us Anything*, treats this as the central mechanism by which our claims have avoided cost. Ostrom's literature shows the structure most clearly, because there the researchers themselves named it and printed it in a journal.

The third finding: Ostrom foreclosed the blueprint reading herself

The third is Ostrom's own. A note in a 1992 working paper she wrote with Michael McGinnis reads: "We do not think it is possible to elucidate necessary AND sufficient principles for enduring institutions as it takes a fundamental willingness of the individuals involved to make any institution work."[16] The two authors foreclosed the blueprint reading themselves, eighteen years before the meta-analyses reached the same conclusion, and Ostrom spent her last decade arguing against panaceas and for a diagnostic approach.[17] Beer hedged the status of his model more than once: "There is no 'correct' interpretation of the VSM," and the choice of five systems is "somewhat arbitrary."[18] He also called it "a cybernetic model of the necessary and sufficient conditions of any viable system,"[19] which is the claim Ostrom refused to make. In the texts we know, nobody has withdrawn it.

A disclaimer is no virtue in itself. This one was load-bearing. It set the expectation against which the later reviews were read, so that a moderate

[15] Beer, *Think Before You Think*, 148; Beer, *Heart of Enterprise*, 275; see also 374 ("there are always pre-symptoms") and 432.

[16] Michael McGinnis and Elinor Ostrom, "Design Principles for Local and Global Commons," working paper D92-6 (Workshop in Political Theory and Policy Analysis, Indiana University, 1992), 40, note 3, https://hdl.handle.net/10535/5460, prepared for the conference "Linking Local and Global Commons," Harvard Center for International Affairs, 23–25 April 1992. The note continues: "No set of logical conditions are sufficient to insure that all sets of individuals are willing and able to make an institution characterized by such conditions work." Earlier drafts of this chapter looked for the sentence in Ostrom's "Institutional Analysis, Design Principles and Threats to Sustainable Community Governance and Management of Commons" (in R. S. Pomeroy, ed., *Community Management and Common Property of Coastal Fisheries in Asia and the Pacific*, ICLARM, 1994, 34–50), which does not contain it.

[17] Elinor Ostrom, "A Diagnostic Approach for Going beyond Panaceas," *Proceedings of the National Academy of Sciences* 104, no. 39 (2007): 15181–87, https://doi.org/10.1073/pnas.0702288104; Elinor Ostrom, "Beyond Markets and States: Polycentric Governance of Complex Economic Systems," *American Economic Review* 100, no. 3 (2010): 641–72, https://doi.org/10.1257/aer.100.3.641.

[18] Beer, *Diagnosing the System*, 144; Beer, *Brain of the Firm*, 98.

[19] Beer, *Heart of Enterprise*, 235.

result could be reported as a moderate result rather than as an embarrassment. We have no such statement on record, and the absence has a cost. Because we have never set out in print what the model cannot do, every finding that the model did not do it arrives as an attack rather than as information.

Where the comparison does not transfer

Common-pool resource institutions are not firms. Membership is frequently inherited and tied to a place, exit is expensive or impossible, and the resource is a physical stock with a regeneration rate that sets the time constants of the whole system. There is no capital structure, no employment hierarchy, and often no competitor. The governance problem is appropriation from a shared stock under temptation to overuse, which is not the problem a manufacturer or a hospital or a software company is solving. And no commercial organization is measured against their success criterion: the institution and the resource are both still there in a century. Ostrom herself later took part in an attempt to carry the principles across. With David Sloan Wilson and Michael Cox she argued that they follow from more general evolutionary principles and apply to a wider range of groups than those that manage common-pool resources.[20] The differences just listed are the ones such a generalization has to survive in a firm, and Chapter 4 asks the same question of the fourth principle.

So the support ratios cannot be carried across as though they were measurements of the VSM, least of all when they happen to suit us. The argument cuts both ways, and the symmetry is the honest part. If the weakness of nesting in Ostrom's sample is only partial evidence against recursion in firms, then the presence of nesting in her table is only partial evidence for it, and the independent-arrival argument in Chapter 10 is correspondingly weaker than the confident form in which we have often made it. A claim cannot be strong corroboration on the way in and weak analogy on the way out. The gap between nesting and recursion narrows the argument once more in both directions, since nesting bears on recursion only where the nested units are viable.

What does transfer is the method: a published instrument, cases coded against it before the outcome is known, independent recoding by people free to disagree, and the disconfirmations printed alongside the rest. None of that is specific to irrigation. It is available to any community with archives and the willingness to open them.

[20] David Sloan Wilson, Elinor Ostrom, and Michael E. Cox, "Generalizing the Core Design Principles for the Efficacy of Groups," *Journal of Economic Behavior & Organization* 90 (2013): S21–S32, https://doi.org/10.1016/j.jebo.2012.12.010. Summarized from the published abstract; we have not read the full article.

The comparison that can be made

The comparison that *can* be made precisely is with Ostrom, and it is worse. On her side stand more than a thousand collected case studies, a four-hundred-and-fifty-five-question coding instrument, eighty-six systematically coded cases, three independent review and reanalysis studies covering ninety-one, sixty-nine and fifty-four cases respectively, and a research program that published its own disconfirmations. On ours, a 2025 systematic review retained twenty-one studies from seven hundred and sixty-eight Scopus records.[21] Eighty-six coded cases against twenty-one studies is a factor of four, so the gap is not really one of counting. It is that one side coded its cases against a published instrument, had the result independently tested three times, and printed the findings that went against it, and the other side has done none of those things.

Part of the difference is resources, and it would be unfair to ourselves to ignore that. Ostrom had a research workshop, doctoral students, decades of funding and a discipline that rewarded the coding of cases. We have consultants with client confidentiality agreements and a small academic literature with no dedicated institution behind it. That explains the difference in counts. It does not explain the three omissions that would have cost us almost nothing: writing down what a successful diagnosis would look like before making one, letting a second person recode a case, and publishing an engagement that did not work.

What this chapter hands forward

Build the register. Ostrom's instrument is published and could be adapted rather than invented. A registry of VSM interventions with prestated success criteria, published boundary judgments and reported failures is the single highest-value program available to this community. It requires no theoretical advance whatsoever, only the willingness to be seen losing.

Priced pilot · Adapt the CPR manual to ten accessible VSM engagement archives. Two researchers can define the codebook, double-code the cases, adjudicate disagreements and publish the pilot in approximately thirty working days, excluding archive permissions. A negative result — diagnoses do not predict the stated outcomes, or cannot be coded reliably — counts against the model and must be published. Chapter 29, *The Two Cheapest Decisive Experiments*, builds the register of interventions and the prediction test on it. This chapter prices only the retrospective pilot. An optional second arm: two coders apply Beer's criterion of viability, "able to maintain

[21] Mohammad Reza Fathi, Soraya Birami, Alireza Payvar, and SeyedAli Doorafshan, "A Systematic Review of the Viable System Model: Applications, Insights, and Future Directions," *Journal of Systems Thinking in Practice* 4, no. 3 (2025): 109–45, https://doi.org/10.22067/jstinp.2025.91552.1136. The PRISMA flow (fig. 2, p. 114) starts from 768 Scopus records and 45 from manual searches and ends with 21 included articles.

a separate existence," to the nested enterprises in a subsample of the coded CPR cases and report their agreement; its price is still to be set.[22]

And take Ostrom's seventh principle as a constraint on cybernetics itself. No amount of internal viability survives the withdrawal of the right to self-organize. That right is a precondition for any design, and the model places it in the containing recursion without regulating its withdrawal. It belongs with the question of who may draw a boundary, which Chapter 29 keeps separate from what the data support. Our best-known case study is the proof. Asked why his team could not have designed for the pressure from outside Chile, Beer answered that it was "like complaining that man, who is supposed to be an adaptive biological system, cannot adapt to a bullet through the heart."[23] Ostrom's seventh principle turns the same condition into a design question: the right to organize has to be recognized and kept, not only survived. We have spent fifty years telling Chile as a tragedy and have not read that difference as a finding.

[22] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 1, where "viable" is defined with the *Oxford English Dictionary*; the criterion is set out in Chapter 10.

[23] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 346: "I have often been asked why we were not able to stipulate a behaviour which would accommodate that threat. It is like complaining that man, who is supposed to be an adaptive biological system, cannot adapt to a bullet through the heart."

Chapter 15

The Thirty-Year Experiment We Never Looked At

In 1999 an agent-based control system called Production 2000+ is installed on a cylinder-head line at the Mercedes-Benz engine plant in Untertürkheim. Its architect is Stefan Bussmann. Behind it stand a reference architecture, an international consortium and a standards track, and in front of it a literature that will accumulate something like fourteen hundred citations.

It passes its performance test. It is deployed for day-to-day operation. It works.

Bussmann's own page records what happened next, in a sentence that repays a second reading: "despite its success during the performance test and during the day-to-day operation, no second production system based on these concepts has been deployed until today."

It runs for five years, until the engine it makes parts for goes out of production and the line is torn down. More than a quarter of a century after the installation it is still the only one.[1]

An experiment that ran alongside us

The installation in Untertürkheim is the visible end of a research program that ran for three decades alongside our own, on problems we have always claimed as ours: how to decompose a working system into parts that hold together without a central schedule, how to keep local autonomy from dissolving the whole, and what happens when the coordinating part is taken away. The program has a reference architecture, a conference series, industrial sponsors and a standards track. We did not read it.

Several literatures reached our components without us, and this is the one that went furthest into practice. Its result is also the least comfortable, because it is largely negative and was published by the people who produced

[1] Stefan Bussmann, "Production 2000+," Stefan Bussmann — Software Agents, http://www.stefan-bussmann.de/en/agents/p2000p.html, for the installation, the performance test, the deployment and the sentence quoted; the reference architecture, the consortium, the standards track and the citation count are sourced later in this chapter. For the five years of operation: Stefan Bussmann, "The Deployment of Production 2000+," http://www.stefan-bussmann.de/en/agents/p2pdeployment.html (consulted 29 September 2026), where the prototype produced cylinder heads for the OM 646 engine from the end of 1999 "until 2005" and was then "torn down due to the end of the product life-cycle."

it, under their own names, in their own venues. Nothing that follows criticizes that literature. It did what we did not: it specified the architecture tightly enough to build, built it, ran it in a plant, and then wrote down that the thing did not spread.

Two looser versions of the charge against ourselves would be wrong. Holonic manufacturing did not succeed where the Viable System Model failed; by the ordinary measure of industrial diffusion it did not succeed either. Nor are holonic architectures the VSM under another name. They descend from Koestler and from control engineering, not from Beer, and their concerns are narrower than ours. The charge is this: a neighboring field ran a long, well-documented and partly negative experiment on architectures with our shape, and organizational cybernetics has no recorded engagement with the outcome. In the sources checked for this chapter we found no cybernetic reading of the holonic deployment record at all. Thirty years of evidence about the fate of architectures like ours sat one journal shelf away, and we did not go and look.

The architecture

Chapter 5, *The Ladder Without a Ruler*, introduces PROSA in passing, as a modeling vocabulary that fixes what its own terms may be applied to. Here the account can be given in full.

Holonic manufacturing systems took Koestler's word[2] into factory automation in the nineteen-nineties. In 1998 Van Brussel and colleagues published PROSA, a reference architecture with a small fixed vocabulary of holon types: product, resource, order, and an advisory staff holon whose failure degrades the holarchy to a heterarchy instead of killing it. Beer has no equivalent design commitment; the VSM's metasystem↗ is not optional in the same way.[3]

We tend to admire two properties of that design together and inherit neither. The first is the closed vocabulary. PROSA offers no open-ended list of things a holon might be. It names a small set of types and requires every element of a plant to be assigned to one of them. That constraint is what lets two engineers modeling the same production cell produce descriptions that can be compared. The VSM has a closed set of types too: "There could not be a 'System Six'," Beer wrote, "because the five-fold system is a

[2]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967), where the word *holon* is introduced.

[3]Hendrik Van Brussel, Jo Wyns, Paul Valckenaers, Luc Bongaerts, and Patrick Peeters, "Reference Architecture for Holonic Manufacturing Systems: PROSA," *Computers in Industry* 37, no. 3 (1998): 255–74. The graceful-degradation property is stated most plainly not in the original paper but in Paul Valckenaers, "Perspective on Holonic Manufacturing Systems: PROSA Becomes ARTI," *Computers in Industry* 120 (2020): 103226: "To qualify as a staff holon (e.g. a scheduler), the staff holon must be no more than an advisor. Removal or disabling of a staff holon must not break the execution system composed of order holons, product holons and resource holons. Performance may degrade, performance predictability may suffer, but the heterarchical operating mode remains operational."

closed system, in logic," and he assigned familiar functions to them, market research and corporate planning to System Four, most management accounting to System Two.[4] What our practice lacks is the rule that makes two modelers converge. Beer declared that there is "no 'correct' interpretation of the VSM" and left the assignment to the observer.[5] The second is the specified failure mode. PROSA states in advance what is lost when the advisory element is removed: performance and predictability degrade, and execution continues in a heterarchical mode. The degradation story belongs to the specification; nobody added it afterward as commentary.

Here the difference is narrower than we used to say. Beer derives each subsystem from what fails without it: System Two is "logically necessary" because without it System One "would go into an uncontrollable oscillation"; System Four is "a necessary condition of viability"; "There has to be a System Five."[6] He also describes the degraded mode. Without System Four "the whole metasystem collapses into System Three," and what is left is a decerebrate cat that still lives: "If you prod its leg, it kicks back."[7] So insight is not what separates us, and the failure logic is in the derivation. What Beer did not do is state it as a requirement that a design must meet, with the performance that remains specified in advance. The engineers wrote the degradation into the architecture as a requirement a candidate design must satisfy; we wrote it into our diagnoses as something a consultant may notice. A requirement can be tested against an implementation.

PROSA had about fourteen hundred citations when checked on 12 August 2026.[8] The figure matters for one purpose: it removes obscurity as an excuse. A diligent reader could not reasonably have missed this literature. It was a well-cited reference architecture in a field with its own conferences and standards work, using a word we also use, for a structure we also describe.

The deployment record

And after thirty years the count of holonic architectures that have actually run a production line is close to one.[9]

---

[4]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 130; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 231 and 179–81; for the six vertical channels, 142–43.

[5]Beer, *Diagnosing the System*, 144; Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 243.

[6]Beer, *Heart of Enterprise*, 177, 143 and 259.

[7]Beer, *Heart of Enterprise*, 265; Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150, where the cat responds "*reactively*" and is "incapable of planning and foresight (Four) and will and judgment (Five)."

[8]Citation counts on 12 August 2026: 1,468 in Semantic Scholar and 1,483 in SciSpace. Counts differ by index and drift upward; "some fourteen hundred" is the honest form of the number and should not be quoted to more precision than that.

[9]This count is ours and it is a count of a negative. We arrived at it by taking the three sources that would report a deployment if one existed — the architect's own record of Production 2000+, the vendor compendium from Rockwell Automation, and the 2020 review of holonic

The genuine case is Production 2000+. Its first prototype was installed in 1999 at the Mercedes-Benz engine plant in Stuttgart-Untertürkheim, controlling workpiece transport for cylinder-head manufacture, and it did run in day-to-day operation; its own architect records that no second production system built on those concepts has been deployed since.[10] The vendor literature is candid. Rockwell's own compendium describes a steel-mill application used only to recommend configurations to an operator, a naval system tested on a scaled demonstrator, and a packing cell that was a university laboratory test bed.[11] A 2020 review of holonic control for Industry 4.0 lists a handful of cases and devotes its conclusions entirely to research directions still to be pursued.[12]

Each of those is real engineering, and calling them non-deployments passes no judgment on their quality. They are simply not production installations, and that distinction matters when the question is whether an architecture reaches the floor and stays there.

---

control architectures — and finding in them exactly one system described as having run in day-to-day production. We know of no systematic census of holonic deployments, and the absence of one is part of the point; a reader who knows of a second production line should treat this sentence as falsifiable and falsify it.

[10] Stefan Bussmann, "Production 2000+," Stefan Bussmann — Software Agents, accessed 2026, http://www.stefan-bussmann.de/en/agents/p2000p.html. Bussmann was the architect of the agent-based control system. His page records that "the first prototype was installed in 1999 at the Mercedes-Benz plant in Untertürkheim," that it "was successfully tested in a performance test and then deployed at the plant for day-to-day operation," and that "despite its success during the performance test and during the day-to-day operation, no second production system based on these concepts has been deployed until today." The primary publication is Stefan Bussmann and Jörg Sieverding, "Holonic Control of an Engine Assembly Plant: An Industrial Evaluation," in *Proceedings of the 2001 IEEE International Conference on Systems, Man and Cybernetics* (IEEE, 2001), 3151–56, which we cite from the review literature rather than from the volume itself.

[11] Pavel Vrba, Pavel Tichý, Vladimír Mařík, Kenwood H. Hall, Raymond J. Staron, Francisco P. Maturana, and Petr Kadera, "Rockwell Automation's Holonic and Multiagent Control Systems Compendium," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 41, no. 1 (2011): 14–30. On the steel application the compendium is explicit: "Although the agent system performed well in all testing scenarios, it was used only for recommending suitable configurations to the operator." The naval work was deployed and tested on the US Navy's Reduced Scale Advanced Demonstrator in Philadelphia, a reconfigurable test platform mimicking the chilled-water system of a DDG-51 destroyer. The packing cell was a physical test bed at the Distributed Information and Automation Laboratory in Cambridge — "a physical model of the production line for packing of customized gift boxes" — and this book previously described it as never physically deployed, which was wrong; it was a university laboratory rig, not a production installation, and that is what the sentence now says.

[12] William Derigent, Olivier Cardin, and Damien Trentesaux, "Industry 4.0: Contributions of Holonic Manufacturing Control Architectures and Future Challenges," *Journal of Intelligent Manufacturing* 32, no. 7 (2021): 1797–1818, published online 22 January 2020, which is why the text calls it a 2020 review. Its industrial cases are few and cited in passing — early automotive work, railway transportation, radiopharmaceutical production — and its conclusions are a list of research directions for future architectures. It does not itself state that holonic architectures remain largely academic; this book previously said that it did, and the sentence has been corrected. The judgement that the architectures remain largely academic is ours, drawn from the paucity of deployments rather than from any statement in the review.

What the count is, and what it is not

The count of one is ours, and it is a count of a negative. We assembled it from the three sources that would have reported a further deployment had there been one, and we know of no systematic census to check it against. By the standards this book asks of others, that is a weak evidential base, and we would rather say so than let the number harden.

The criterion also does work that should be visible. We counted systems described as running in day-to-day production, which is demanding. A looser criterion that counted pilots, advisory installations and demonstrators would return a larger number and a different chapter. The claim therefore has a conditional form: on the strictest reading of deployment, and on the sources we have checked, the record contains one case, and the field's own reviews do not point to others.

The architect has given his own reason. The system, Bussmann writes, performed well and "lived up to all its technical promises"; the trouble was economic. Flexibility "is a future advantage which requires an immediate investment," so the business case was weak: "You must be convinced that you require the flexibility - you cannot prove it." Gantry loaders gave comparable productivity gains with little flexibility and at lower cost, and nearly all the systems Daimler installed afterwards were of that type. He adds that the assessment is "strongly biased towards the automotive industry."[13] That is one account of one plant, and other explanations remain open. The installed base of conventional control systems is large, and whoever admires the new architecture does not pay the replacement cost. Procurement, maintenance skills and liability all favor the incumbent. The problems holonic control solves may not be the problems these plants were losing money on. The advantage may be real but visible only at disruption rates ordinary plants do not experience. Nothing we have read settles among these beyond the one plant, and a chapter that picked one would be doing what we accuse ourselves of elsewhere in this book: preferring a satisfying account to a tested one.

What the record removes

That is a thirty-year controlled experiment, and cybernetics has not looked at it. Holonic manufacturing had what the VSM never had: assignment rules that make two models of the same plant comparable, an international program, industrial partners, a conference series, a vendor

[13]Stefan Bussmann, "Production 2000+: Lessons Learned," Stefan Bussmann — Software Agents, http://www.stefan-bussmann.de/en/agents/p2plessons.html (consulted 29 September 2026), which draws on Klaus Schild and Stefan Bussmann, "Self-Organization in Manufacturing Operations," *Communications of the ACM* 50, no. 12 (2007): 74–79, https://doi.org/10.1145/1323688.1323698. The gantry-loader sentence reads: "nearly all manufacturing systems that Daimler has installed since the project P2000+ are of this type (at least up to 2003)."

research laboratory and a standards track, and it still did not ship.[14] Whatever explains why cybernetic architectures do not reach production, "we never wrote it down properly" is not it.

The inference is narrow and should stay narrow. It does not show that specification is worthless. A field that had not specified could not have produced the record we are reading, and only formalization made the graceful-degradation property statable. Nor does it show that our own defects, cataloged in Part One, are harmless. It removes a single explanation we have leaned on for years: that the model's limited industrial reach is a writing-up problem, and that a tighter notation or a proper reference architecture would convert interest into installations. A neighboring field did the writing up, to a standard we have never reached, and the installations did not follow.

The transfer is not automatic. Holonic architectures are software control systems for physical plant. The VSM is a model of regulation in any viable system, applied mostly to organizations of people, and a failure to diffuse might have causes in the first case with no analog in the second. The contrast is less clean than it looks, because the VSM's own built case was also a network and software control system. That is why the right response is a comparative study rather than a conclusion, and the receipt below buys one.

The case we built ourselves

The comparison so far leaves out the one case in which a VSM-based architecture was specified, built and run. Chapter 8 tells the story of Project Cybersyn. Here only its record is needed, set beside the holonic one.

Between late 1971 and September 1973 Beer and a Chilean team built a telex network, Cybernet; a suite of statistical programs, Cyberstride, "needed to provide statistical filtration for all homeostatic loops at all levels of recursion"; and an operations room that was "in experimental working

[14]The reference architecture is PROSA, above. For the consortium and the standards track: James H. Christensen, "Holonic Manufacturing Systems: Initial Architecture and Standards Directions," in *Proceedings of the First European Workshop on Holonic Manufacturing Systems* (HMS Consortium, Hannover, 1 December 1994) — the HMS consortium being a test case of the international Intelligent Manufacturing Systems programme; IEC 61499-1:2012, *Function Blocks — Part 1: Architecture*, ed. 2.0 (International Electrotechnical Commission, 2012), the distributed-control standard that came out of that work, alongside IEC 61131-3, *Programmable Controllers — Part 3: Programming Languages*, now in its fourth edition (International Electrotechnical Commission, 2025), which it did not displace; and Alois Zoitl et al., "The Past, Present, and Future of IEC 61499," in *Holonic and Multi-Agent Systems for Manufacturing*, ed. Vladimír Mařík, Valeriy Vyatkin, and Armando W. Colombo, Lecture Notes in Computer Science 4659 (Springer, 2007), 1–14. The conference series is HoloMAS, from 2000, when it began as a workshop at the DEXA conference. The vendor laboratory is the Rockwell Automation Research Center in Prague, whose output is the compendium cited above. Valckenaers, "PROSA Becomes ARTI," records the standard's fate from inside: "IEC-61499 is struggling to achieve critical user mass, which is imperative for a technology of its complexity."

order on 10th January 1973."[15] By Beer's count, "something like seventy percent of the socio-industrial economy was operating this system, involving about four hundred enterprises"; in *The Heart of Enterprise* he gives seventy-five percent at the time of the coup.[16] Both figures are his own, and they are not the same figure. Eden Medina's archival history records a smaller working state: at the end of November 1972 the temporary suite was checking seventy-four production indicators for anomalies, twenty-six were ready for the permanent suite and a hundred and eighty more were in preparation.[17] The network carried the government's distribution of essential goods through the strike of October 1972, on Beer's account and on Medina's, and, on Beer's, through a second strike in August 1973.[18]

The end is documented; the performance is not. After the coup of 11 September 1973 the military stopped work on the project and abandoned or destroyed what the team had built.[19] Beer, asked why the design could not accommodate the pressure from outside Chile, compared the question to complaining that a man "cannot adapt to a bullet through the heart."[20] Five months before the coup he had written his own assessment. The technical objectives had been met, the technocratic ones might yet be, the social outcome was ambiguous and the political objective would not be reached; each member of the team weighted the four differently, which "would explain some current confusion, and the disagreement about success." In the same paper he recorded that the Chilean establishment was "lauding and gladly incorporating" some components, Cyberstride and the operations room among them, and ignoring others, algedonic metering among them.[21] No criterion of success was stated before the system ran, so there is nothing to test its record against. Medina met people who thought the project did not warrant historical attention because it "never reached completion."[22]

Three records now sit side by side. Production 2000+ was specified, built, passed its test, ran, and was not repeated. The sociotechnical designs of Chapter 13 were built many times and did not diffuse. Cybersyn was specified, built in part, ran for less than two years and was ended by force, after its designer had watched it being taken up piecemeal. None of the

[15] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 270. Beer's account of Chile runs from 245 to 348 and is his own, first-person record.

[16] Beer, *Brain of the Firm*, 264; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 297: "we were (at the time of the coup on 11th September 1973) monitoring 75 per-cent of the social economy on a single IBM 360/50 — down to plant level." These two figures, and "Roughly 60% of all enterprises constituting the socio-industrial economy" in *Brain of the Firm*, 311, measure different things at different dates and should not be merged.

[17] Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 167.

[18] Beer, *Brain of the Firm*, 312–14, where "One senior minister said flatly that the government would have collapsed that night if it had not had the cybernetic tools" (314), and 346 for August 1973; Medina, *Cybernetic Revolutionaries*, 141 and 148. Chapter 16 gives Shalizi's reading of the same episode.

[19] Medina, *Cybernetic Revolutionaries*, 211.

[20] Beer, *Brain of the Firm*, 346.

[21] Beer, *Brain of the Firm*, 339–40, quoting his paper "On Decybernation" of April 1973.

[22] Medina, *Cybernetic Revolutionaries*, 10.

three stopped for want of a specification. The causes differ on their face, and whether they share anything beneath it is the question the receipt below prices.

Four more literatures on the same shelf

The holonic case is the largest of the fields we did not read, and the others share its shape: mature work, adjacent to a claim we make, available for decades, unengaged.

Qualitative comparative analysis is the clearest instance. Its foundational statement appeared in 1987, and its central finding is that configurations of conditions can be jointly effective while no single condition is necessary or sufficient on its own. Our five-system claim is stronger than that. Beer holds each subsystem necessary, System Two "logically necessary to any viable system", System Four "a necessary condition of viability", "There has to be a System Five", and the set sufficient: "a cybernetic model of the necessary and sufficient conditions of any viable system."[23] He also holds that the functions are always performed somewhere. System Two "is always present, or the organization would shake itself to pieces", and "System Four exists — it has to: but it is disseminated."[24] A necessity claim of that kind can be tested only once absence has been coded in advance, and the coding in Chapter 28 separates a function performed by a designated role, a function performed unofficially, and a function nobody performs. QCA's necessity and sufficiency analyses are built for exactly such claims. The method has existed throughout the period in which we have been asserting ours, and we have tested ours by illustration.[25]

Measurement shows the same pattern. Chapter 14 describes the published coding manual behind Ostrom's program and prices a ten-archive pilot against it. The pilot is cheap; the expensive path is the alternative. Building a new VSM questionnaire of our own before adapting an instrument already validated on comparable material would repeat, in the present tense, the neglect this chapter documents about the past.

Sociological systems theory is the third. Chapter 16 carries the full citations to Achterbergh and Vriens (2009), Schwaninger and Groesser (2012), and Johnson and Leydesdorff (2015). What matters here is narrow: "the traditions cannot be combined" is no longer an inference available without first reading the attempts, and the attempts exist.

The fourth entry is the cheapest and the least developed. Troncale's system process theory and its linkage propositions offer an explicit way to pair recurring processes with pathologies and channels, which is close to

[23] Beer, *Heart of Enterprise*, 177, 143, 259 and 235.

[24] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 66; Beer, *Heart of Enterprise*, 143, where System Four "lacks integration"; see also 265 on the "decerebrate cat".

[25] Charles C. Ragin, *The Comparative Method: Moving beyond Qualitative and Quantitative Strategies* (University of California Press, 1987).

something we need.[26] The negative search is itself dated: the material was absent from every manuscript draft checked on 12 August 2026.

Receipt — a three-way comparative post-mortem. Compare three records: holonic manufacturing (Bussmann's account, the Rockwell compendium and the Derigent review, followed by targeted checks for a second production deployment), the sociotechnical designs of Chapter 13, and Cybersyn (Beer's account in *Brain of the Firm* and Medina's archival history). Budget twelve researcher-days for the holonic record and about six more for the other two. Pre-register rival explanations for non-diffusion, including Bussmann's weak business case for Production 2000+; Beer's own for Cybersyn, the piecemeal incorporation of components by an existing establishment; and Beer's own general explanations, paradigm closure, which he called the "Noah's Ark Syndrome," and the view that the battle for Cybersyn's reputation had been "lost in advance."[27] Record for each case which explanations its sources can test and which they cannot. If one explanation fits all three, retire "write it down better" as the remedy; if none does, state the differentiating mechanism. The holonic count is explicitly falsifiable and must be rechecked before print.

[26]The early statement is Len Troncale, "Linkage Propositions between Fifty Principal Systems Concepts," in *Applied General Systems Research*, ed. George J. Klir (Plenum, 1978), 29–52, https://doi.org/10.1007/978-1-4757-0555-3_2, which we have not read. The later statement we have read is Len Troncale, "Systems Processes and Pathologies: Creating an Integrated Framework for Systems Science," paper for the INCOSE International Symposium (IS13), 2013: fifty-five systems processes "joined by more than 200 Linkage Propositions," and a systems pathology derived from them, with classes such as dysfunctions of feedback and dysfunctions of flows.

[27]Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 183–84, where paradigms that "RESIST CHANGE" are "the Noah's Ark Syndrome"; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 336: "the battle of Ashby's Law for the reputation of Project Cybersyn had (as was said earlier) been lost in advance."

Chapter 16

The Objections That Land, and What They Cost

In 1981 Werner Ulrich publishes twenty-seven pages in the *Journal of Applied Systems Analysis*, putting to Project Cybersyn the questions C. West Churchman had formulated for any systems study. His conclusion is one sentence: "Cybersyn's alleged goal of devolving power to workers is not confirmed by its design." Beer sent his comments on tape, and the tape, if it survives, has not been found.

Apart from five pages Beer printed two years later, which we have not yet read, that is the whole of this field's reply. In the same year the second edition of *Brain of the Firm* addressed the charge of technocracy in general terms, conceding that the participatory design "was not always being done in the intended spirit"; it does not mention Ulrich and does not answer his argument about the client.[1] The objections in the pages that follow deserved a better reply.[2]

What lands, and by what standard

The chapters before this one kept us company: neighboring literatures that arrived at our components without us, mostly without knowing we had them. What follows is opposition, and a reviewer will turn here first.

Three objections land: Werner Ulrich's, F. A. Hayek's and Robert Cooper's. We take them first, in that order, before the positions that do not land, and they turn out to be one wound seen from three directions. It is a wound about boundaries.

The standard is narrow, and it has to be stated, because a chapter of this kind can reach any verdict by adjusting what counts as a hit. An objection lands here when it identifies something the Viable System Model↗[3] needs

[1] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 271, where the lapse "seemed likely" to have "aided detractors who later wished to call Cybersyn 'technocratic' "; see also 322 and 328.

[2] Werner Ulrich, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59 — twenty-seven pages. The quotation and the taped reply are from the paper and its acknowledgements; the note below on Ulrich sets out the argument in full and what its standing rests on. The five categories are C. West Churchman's, applied by Ulrich.

[3] Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

in order to do the job it claims, cannot supply from its own resources, and has not been given by any amendment since. An objection that prices an amendment (an added index, a demoted claim, a stated scope condition) does not land in that sense, however uncomfortable it is to receive. We want to tell a refutation from a repair bill, and to pay the repair bill rather than argue with it.

In one respect all the positions collected here are of the more serious kind. None of them thinks the model merely imprecise or under-tested. They think the job it proposes is not well posed: there is no bounded steerable object to diagnose, or no vantage point from which to diagnose it, or no innocent way to draw its edges. If any of them is right in that strong form, better measurement makes things worse, because it refines an instrument aimed at nothing.

One structural fact belongs before the detail. It is the strongest defensive resource this book has, and it is a fact, not a rhetorical move: these positions are mutually inconsistent. Cooper and Burrell classify Luhmann among the modernists.[4] DeLanda is a realist about emergent wholes at several scales, which actor-network theory denies; actor-network theory abolishes the levels DeLanda keeps. Ulrich and Jackson both presuppose that a systems study is a legitimate thing to do properly, which the process theorists deny. The model does not have to survive all of them, only whichever of them survive each other.

A note on language

This chapter borrows four vocabularies, and each of them uses a familiar word in an unfamiliar way.

Luhmann's *communication* is not people talking: it is the unit a social system is made of, and it excludes the people. His *operational closure* does not mean isolation; it means that no operation begins in one system and completes in another. DeLanda's *relations of exteriority* is the claim that a component is not constituted by the role it plays in a whole — which is why he can detach it and the model, as we teach it, cannot. Ulrich's *boundary judgment* is not a boundary but the decision, usually unnoticed and always somebody's, about where one goes. And Jackson's *coercive* classifies a context in which interests cannot be reconciled; it is not an accusation against anyone standing in it.

Where these authors and this book use the same word for different things, we have kept theirs and marked ours.

After the three that land come four positions that do not land in the form they are usually given: Luhmann's operational closure, the assemblage and network arguments about levels, Michael Jackson's placement of the model within a system of methodologies, and the wider process argument of which Cooper's boundary point is one part. For each we say what it does

[4]Robert Cooper and Gibson Burrell, "Modernism, Postmodernism and Organizational Analysis: An Introduction," *Organization Studies* 9, no. 1 (1988): 91–112. Luhmann appears there alongside Bell and Habermas as an exemplar of modernist discourse, which is the classification the text relies on.

establish. Each converts into something we owe, and the closing section prices the only one of them that can be paid immediately.

Ulrich: whose boundary, and who witnesses

Werner Ulrich is the sharpest of these critics because he does not deny that systems modeling works. He denies that it can be innocent.

Nor is he a distant critic. His 1981 case study, built on Beer's own reports and checked by an insider,[5] finds that the design does not do what its designer said it did. The client Beer declared was the workers; the client the design serves is the government, which alone can act on the system's reports without a homeostat↗ above it, and an algedonic↗ signal is a cry of pain and also a transfer of control upward. Hence the sentence in the paper that is hardest to answer: "what distinguishes Cybersyn from a traditional state bureaucracy is not a less hierarchical structure but the greater

[5]Werner Ulrich, "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59. A Spanish translation appeared as *Una crítica de la razón cibernética pura (el caso del proyecto CYBERSYN)*, Memos de Investigación 36 (Universidad de los Andes, 1989). Earlier drafts of this book recorded that we had not been able to obtain the text and were relying on Ulrich's own one-line description of it; we have read it since, and this note replaces that admission. The paper has two halves. The first argues that systems theory has been absorbed by a "scientistic" school which mistakes theoretical for practical reason, and proposes a critically-normative paradigm against what Ulrich names the *syntactic fallacy* — taking Ashby's law as a criterion of systems rationality when the level of organization is socio-cultural rather than organic. The second applies five of Churchman's categories — client, purpose, decision maker, expert, guarantor — to Cybersyn, and concludes that "Cybersyn's alleged goal of devolving power to workers is not confirmed by its design. Instead, its strongly hierarchical organisation and its concept of 'autonomy' one-sidedly serve the top decision maker, the government." Three things should be said about its standing. Ulrich states at the outset that the political purpose Cybersyn served is *not* the subject of his critique — "It is not my intent to question the social and democratic ideals that both Salvador Allende and Stafford Beer associated with the project" — and closes by regretting that his argument falls on "one of the very rare efforts undertaken by systems scientists in the service of democratic values." The account was revised after comments from Fernando Flores, Allende's Secretary of Economics, and its accuracy was confirmed by "an anonymous reviewer who was involved in the project"; where Ulrich contradicts the contemporary *New Scientist* coverage, he does so on Flores's authority. And on the question of a reply: the acknowledgements record that "Stafford Beer too was kind enough to transmit his oral comments on tape to me, though he could not agree with my critique." Beer did answer in print: Stafford Beer, "A Reply to Ulrich's 'Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics,'" *Journal of Applied Systems Analysis* 10 (1983): 115–19, listed and quoted in M. C. Jackson, "An Appreciation of Stafford Beer's 'Viable System' Viewpoint on Managerial Practice," *Journal of Management Studies* 25, no. 6 (1988): 557–73, at 572 and 564; we have not read it, and earlier drafts of this book said no reply had been published. Beer's letter in *New Scientist* of 22 March 1973 answers Herbert Grosch's letter of 15 March, not Ulrich (Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 186–87 and 289 n. 55), and the tape, if it survives, would be worth finding; the search is priced once, in the receipt of Chapter 8. One limit of the paper is worth stating in its favor as well as against it: it is built on *Decision and Control*, *Brain of the Firm*, *Platform for Change*, the Fanfare lecture (the Richard Goodman Memorial Lecture, Brighton Polytechnic, 14 February 1973, reprinted in *Platform for Change*, 421–52) and the 1974 Zaheer lecture, and engages neither *The Heart of Enterprise* nor *Designing Freedom*, so it does not meet the Beer of Chapter 26, *Freedom as a Quantity*.

speed at which vertical communication (and control) takes place." That is the objection Part Four is built to answer, made in 1981 and aimed at the founder.

Beer's own pages answer part of this and concede part. The design was meant to be the reverse of a transfer of control: "the schema represents maximum decentralization, since any given level of recursion receives directly only its own Cyberstride reports", and the algedonic signal was reserved for a level that failed to deal with its own trouble, so that "if all concerned do their agreed jobs properly, algedonic signals will rarely be fired."[6] He had also conceded, as the opening of this chapter records, that the participatory design was not always kept, and he reported that when the Minister drew back from Cybersyn in December 1972 the gratification "was experienced by the technocratically minded, who wanted effective management regardless of the political framework."[7] Ulrich's point survives as a claim about who sets the filters and the thresholds, a question these passages do not settle.

Chapter 8, *The Loop That Stayed an Intention*, narrates the episode, including the taped comments Beer sent. Beer also published a short reply in the same journal in 1983, which we have not yet read. Here the story matters less than the standing of the argument, and that is undiminished. The paper was revised in the light of a participant's comments and checked by another. It makes no claim about the political purpose the project served, and says so twice. It was aimed at a designer who was alive, working and perfectly capable of answering it. Beyond Beer's own five pages, the reply was ours to write, and we did not write it.

The critique generalized into a method. Critical systems heuristics asks twelve boundary questions, each in an *is* and an *ought* mode: who is and ought to be the client; what is and ought to be the purpose; what is and ought to be the measure of improvement; who decides and what they control; who counts as an expert and what counts as relevant knowledge; what is assumed to be the guarantor of success; and, at the end, who is witness to the interests of those affected but not involved, and what secures their emancipation from the premises of those who are.[8] Chapter 4, *Channels Without Owners*, sets those questions beside the channels the model does specify. Here we ask them of the diagnostic act itself, where they were always pointed.

Applied to us, they are uncomfortable in a precise way. The model leaves the naming of identity to the eye of the beholder and assigns the naming

[6] Beer, *Brain of the Firm*, 275 and 276.

[7] Beer, *Brain of the Firm*, 322.

[8] Werner Ulrich, *Critical Heuristics of Social Planning: A New Approach to Practical Philosophy* (Paul Haupt, 1983); Werner Ulrich and Martin Reynolds, "Critical Systems Heuristics," in *Systems Approaches to Managing Change: A Practical Guide*, ed. Martin Reynolds and Sue Holwell (Springer, 2010), 243–92. The ordering used here is Ulrich's own: the twelve questions fall in four groups of three — sources of motivation, of power, of knowledge, of legitimation — and each is asked in an *is* and an *ought* mode. The measure of improvement is question 3, the guarantor question 9, and the witness to those affected but not involved question 10.

only loosely. Purpose is imputed by "a GOOD OBSERVER" and stated for the system by the higher recursion; the System-in-focus is chosen by the diagnostician; Cybersyn's measures were to be agreed with the people in each plant.[9] The model does not say who is *entitled* to do any of this, how it is recorded, or how a disagreement is settled. It does not say who fixes the recursion levels↗, or on whose authority. It does not say whose measure of improvement the System One↗ indices encode, though they always encode somebody's, and Beer's own Chile account concedes that the participatory design was not always kept. Its assumed guarantor of success is either viability itself or the model's own theorems. Both are internal to the designer's frame, so neither can justify anything to someone standing outside it. Beer's own stated guarantor is predictive utility, which could lie outside the frame; the tradition has not specified the predictions that would put it there.

The tenth question is this chapter's line. The algedonic channel has no seat for the person who is affected by an operation and was never a member of it. Pain is signaled by the units that feel it from inside the system. The neighbor downstream of a plant is not a unit, has no channel, and appears in the design only as environment.

Now what the objection establishes, and what it does not. It does not establish incompatibility. Nothing in critical systems heuristics contradicts the Viable System Model; the two operate at different logical levels, and Ulrich interrogates the boundary that Beer designs inside. Nor does it establish that the model conceals boundary judgments deliberately, and Ulrich does not say so. It establishes that a diagnosis produced with our tools does not justify itself, and that the missing justification takes the form of answerable questions with names attached. System Five↗ and the algedonic channel are, in intent, a witness function. The model has the slot. Beer states who ought to be there, "All stakeholders are represented in System Five," with a consumers' council among his examples,[10] but he gives no procedure by which those affected and not involved obtain or use the seat, and we have not supplied one. That defect can be fixed, and fixing it is the cheapest high-value reform available to this tradition. The receipt at the end of the chapter buys exactly this and nothing else.

Hayek: the knowledge that cannot be sent

Hayek put the second objection in one sentence: "The sort of knowledge with which I have been concerned is knowledge of the kind which by its

[9] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 99 and 16; Beer, *Brain of the Firm*, 271. Purpose is "a subjective phenomenon, rather than a property of the system independent of its instigators, participants, and observers": Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 158. On predictive utility, Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113 and 243.

[10] Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 149; Beer, *Heart of Enterprise*, 263–64.

nature cannot enter into statistics and therefore cannot be conveyed to any central authority in statistical form."[11] It has never been answered, in this tradition or, so far as we can find, anywhere adjacent to it.

The objection comes in layers, and treating it as one is how we keep failing to meet it. Three layers are Hayek's own. The last two belong to his modern defenders, and Cosma Shalizi has put them most sharply.[12] First, the knowledge is tacit, constituted in the act of using it, not merely dispersed. Second, it is *discovered* rather than transmitted: the data a regulator would need does not exist until rivalrous action generates it.[13] Third, there is a measure problem. Indices must be chosen in advance. Beer's actuality-capability-potentiality triple is fixed beforehand, by the analysis teams and, by design if not always in practice, the plant's workers, and Beer bounds potentiality by "what is already known to be feasible,"[14] while Hayek's circumstances of time and place are by definition the ones the modeler did not anticipate. Fourth, there is an objective-function problem. Even with perfect information, weighting goods against each other requires preferences that differ between people and that no mechanism reconciles satisfactorily. In the VSM that weighting has to be smuggled into System Five, which Beer equipped with deliberative devices, Team Syntegrity and the algedonic meter, and with no procedure that reconciles divergent preferences and could claim legitimacy for the result.[15] Last comes the computational layer, the weakest of the five and the one most often mistaken for the whole. When our literature answers the objection at all, it answers this layer, because it is the only one that hardware improvements can address.

Note where the third layer lands. Hayek's measure problem and Ulrich's third boundary question are the same question, asked by an economist and a philosopher who would agree on nothing else: whose measure of improvement is this, and what secures it against the improvements it cannot register? When two hostile literatures converge on a point, the point is load-bearing.

Beer shared Hayek's premise. He wrote that the prerogative to intervene indiscriminately "does not have Requisite Variety," and he answered it with

[11] F. A. Hayek, "The Use of Knowledge in Society," *American Economic Review* 35, no. 4 (1945): 519–30. The sentence quoted is on 524, in the last paragraph of section IV. The phrase "the particular circumstances of time and place" first appears on 521 and recurs on 522 and 524.

[12] Cosma Shalizi, "In Soviet Union, Optimization Problem Solves You," *Crooked Timber* (blog), May 30, 2012. The objective-function layer is his: planners cannot take the relative weighting of goods as given, because "people have different preferences" and "there is no universally satisfactory mechanism for making such choices." The computational layer is the arithmetic of the same essay.

[13] F. A. Hayek, "Competition as a Discovery Procedure," trans. Marcellus S. Snow, *Quarterly Journal of Austrian Economics* 5, no. 3 (2002): 9–23, a translation of a lecture first published in German in 1968. This is the source for the layer that treats the relevant data as generated by rivalrous action rather than merely dispersed; the 1945 paper does not carry that argument on its own.

[14] Beer, *Brain of the Firm*, 163 and 271.

[15] Beer, *Beyond Dispute*, 149–59; Beer, *Brain of the Firm*, 283.

autonomy rather than with central knowledge.[16] The best defense available to us is that Cybersyn was not built to compute a plan. It attenuated↗, and in that it was a Hayekian architecture, leaving the decision to the person on the spot and escalating only the residue. Chapter 8 sets out this reading of the design, and we still think it is the right reading of the intent.

The historical record neither confirms nor refutes the defense. By Beer's own count the system ran across some four hundred enterprises, but never under conditions that could test the claims at issue: the operations room was only "in experimental working order" in January 1973.[17] Cosma Shalizi's empirical verdict, that to call Cybersyn a decentralized planning system "is nonsense" and that at best it would have been a dashboard, is set out in Chapter 8.[18] Beer's first-person account reads October 1972 as filtration, with some signals algedonic and others attenuated, and he staged and rejected the objection that the computers were "merely frills."[19] We have no evidential answer to either reading, because the trial ended before it could produce evidence either way.

So the measure-problem layer stands unanswered. Preselected indices cannot certify that they captured the surprises nobody selected for. No amount of real-time reporting closes that gap, because the gap lies in the choice of what to report, and the choice is made before the surprise arrives. We do not have a reply to this. We have a research program that would begin to constitute one, and Part Five prices its first two steps.

One counter-argument has never been properly staged, and staging it falls to us. Hayek's price system is itself a variety attenuator, carrying a single scalar per commodity. The Ashbyan question is whether one scalar has requisite variety↗ for every regulatory task, and whether its settling time is short enough for every disturbance. Beer came close. He argued that a money-supply regulator lacks requisite variety for an economy that learns,

[16] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 38: "the so-called prerogative to intervene indiscriminately does not have Requisite Variety. It cannot be competently done."

[17] Beer, *Brain of the Firm*, 264 and 270; at 263, "nothing was known about the performance" of the forecasting filter in advance.

[18] Cosma Shalizi, "Project Cybersyn," Notebooks, http://bactra.org/notebooks/cybersyn.html, consulted 29 September 2026; the page is revised continuously and then showed "Last update: 01 Aug 2026" (first version 30 January 2022). The quoted verdict is that "to describe it as a decentralized planning system is nonsense"; the supporting points are that there was "the one (!) mainframe computer available to the project," that the lag between measurement, processing and delivery "was so long that it seems to have been a complete flop," that, in words Shalizi quotes from Eden Medina, "dismantling one of the primary safeguards of [enterprise] autonomy might have been as easy as having someone from the telex room walk down the hall," and that what the project would have produced at best was a "pioneering example of what we'd now call a 'dashboard.'" The telex-room sentence is Medina's, not Shalizi's: Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 184, where it reads "of that autonomy" (Shalizi cites the passage as 183–84). For the October 1972 strike and the scale at which the system actually ran, Medina, *Cybernetic Revolutionaries*, chap. 5, "The October Strike," 141–70.

[19] Beer, *Brain of the Firm*, 312–13 and 394. The account is declared first-person and partisan at 245–46.

and called money "the metric of *constraint*."[20] Nobody, Beer included, has applied the argument carefully to Hayek's price mechanism itself. It is an opening, not yet a citation, and it is the sort of work this book exists to commission.

Cooper: the act that manufactures the object

The third objection that lands is the least often stated and the hardest to shake off. It belongs to a family of process arguments, most of which do not land, so the one that does has to be separated from the rest.

"Change must not be thought of as a property of organization," write Tsoukas and Chia. "Rather, organization must be understood as an emergent property of change."[21] The family position is that an entity is an artifact of a description, and that the description moves more slowly than what it describes. Applied here: by the time the System Ones have been identified, the boundaries drawn and the indices chosen, the flux has reorganized. A diagnosis produces a system-effect rather than a snapshot of a system. Then it feeds back, because people rearrange themselves to fit the model, and the model's apparent accuracy becomes a performative achievement rather than a finding.

Robert Cooper adds the edge that does the damage. In our compression of his argument, the diagnostician's first act, drawing the boundary, is the act that manufactures the object, and it is exactly the act for which the method supplies no procedure of examination.[22] His complaint against systems theory is that it most neglects the concept which originates system: boundary, division.

Beer named the act. "A system is not something given in nature, but something defined by intelligence"; two observers "will not necessarily agree on … the boundaries"; the analyst should "not gaily accept the system delineated by the conventions of management structure." His workbook makes the boundary its first task.[23] He even revised one of his own first

[20] Beer, *Brain of the Firm*, 308; Stafford Beer, *Platform for Change* (John Wiley & Sons, 1975), 159. See also *Decision and Control*, 139.

[21] Haridimos Tsoukas and Robert Chia, "On Organizational Becoming: Rethinking Organizational Change," *Organization Science* 13, no. 5 (2002): 567–82, at 570. Tsoukas and Chia are frequently enrolled as Whiteheadians and are not: their process sources are James, Bergson, Mead and Rescher. The Whiteheadian strand in organization studies runs through Robert Cooper and later writers. Chapter 10, *Six Components*, names Whitehead for the position, which is right for the position and wrong for these authors.

[22] Robert Cooper, "Organization/Disorganization," *Social Science Information* 25, no. 2 (1986): 299–335. Cooper's complaint against systems theory is that "the very concept which originates system — namely, boundary or division — is the one which is most neglected." We have that sentence from the article's own notes; we did not have access to the full body of the text, and the sharper formulation attributed to him here — that drawing the boundary is the act that manufactures the object — is our compression of the argument rather than a quotation. See also Cooper and Burrell, "Modernism, Postmodernism and Organizational Analysis."

[23] Beer, *Decision and Control*, 242, 243 and 246; Beer, *Diagnosing the System*, 2 and 16.

cuts in print, the Chilean model that had mixed two levels of recursion, and what triggered the revision was events, not a method.[24]

Against the general form of the argument Beer built a defense in advance, and it is a real one. The Three-Four homeostat is the response to a description that lags its object. System Four↗'s business is continuous re-modeling, and the insistence on real-time data and short settling times is an attack on lag as such. Process philosophy names a problem the model was partly designed against. It also offers no criterion for telling a better abstraction from a worse one, and some abstractions kill people when they fail. Whitehead's fallacy, in any case, was mistaking the abstract for the concrete, not abstracting.[25]

Cooper's version survives all of that, because it concerns the first cut rather than lag. Nothing in the Three-Four homeostat examines the act that produced the object the homeostat then tracks, and no amount of faster re-modeling reaches back to it. Part One has already made this concession in its own terms, and the convergence is now visible. Ulrich says the first cut is somebody's and should be published with a name on it. Hayek says the first cut fixes what can afterward be measured. Cooper says the first cut makes the object. Three literatures with nothing else in common have found the same act at the start of our method: named by Beer as the observer's, revised once in public, and never given a procedure.

Luhmann: nothing crosses the boundary

The remaining positions are serious and do not land in the form they are usually given. Luhmann's is the most metaphysically ambitious.

Niklas Luhmann's social systems consist of communications, not of people or machines or actions. Psychic systems are separately closed and belong to the environment. Organizations reproduce themselves out of *decisions*, communications that announce their own contingency, that they could have been otherwise. Their structures are decision premises: programs, channels, personnel, and the undecidable premises we call culture.[26] Operational closure carries the weight. No operation begins in one system and completes in another. Information is not transported; the system produces it internally, by its own distinctions.

Follow that through the VSM and the damage is specific. On this account a System One operational unit (people, plant, process, making the thing the

[24] Beer, *Brain of the Firm*, 322–26.

[25] Alfred North Whitehead, *Science and the Modern World* (Macmillan, 1925), where the fallacy of misplaced concreteness is named, and *Process and Reality: An Essay in Cosmology* (Macmillan, 1929), where it is restated. The fallacy is mistaking an abstraction for the concrete fact it abstracts from, which is not an objection to abstracting.

[26] Niklas Luhmann, *Social Systems*, trans. John Bednarz Jr. with Dirk Baecker (Stanford University Press, 1995); Niklas Luhmann, *Organization and Decision*, ed. Dirk Baecker, trans. Rhodes Barrett (Cambridge University Press, 2018). The account of organizations as reproducing themselves out of decisions, and of decision premises as programs, communication channels and personnel, is from the latter.

enterprise exists to make) is not a unit of the social system at all. System Three↗ presupposes a position from which the operations it regulates can be observed. But management is itself decision-making, one more run of communications inside the same autopoiesis, not a stratum above it. "System Three regulates System One" is a self-description the organization produces about itself, and the model mistakes the self-description for the structure.

System Four takes the sharper damage. Environmental scanning does not deliver the environment; it produces internal communications about a distinction the system itself drew. So the celebrated homeostat between outside-and-then and inside-and-now couples no internal reality to an external one. It is an internal oscillation between two self-produced descriptions, with no channel by which the environment could correct either. Beer's requisite-variety justification for System Four is stated in *The Heart of Enterprise* as if it assumed exactly the access closure denies, although by 1980 Beer had accepted that "we observers invent the codes" and had defined closure as "a *self-referential process*, and not the isolation of the system within an adiabatic shell."[27] And System Five, on this reading, does not close the system. It is where the system de-paradoxifies its own founding paradox, that decisions are only needed where the matter is undecidable: the fiction by which it conceals its contingency.

The reply has four parts. The first is that "cannot be steered" is a strong reading, and Luhmannians dispute it among themselves. Helmut Willke built an entire steering theory inside the tradition on *Kontextsteuerung*: you cannot steer a closed system directly, but you can alter the contexts and couplings within which it steers itself. That is a description of what System Three actually does. System Three commands little. Besides the command channel, on which it "exerts authority",[28] it sets the resource bargain, the accountability and the standards, the context within which System One regulates itself. Beer's principle of autonomy and Willke's context steering are the same idea, arrived at from opposite ends.[29]

The second is that closure permits a great deal. Structural coupling means systems can be durably linked so that each is a standing source of irritation for the other; contracts, budgets, membership and technology are structural couplings. Second-order observation, observing observers and one's own self-descriptions, can be institutionalized as routine, which is functionally close to what we call Systems Four and Five. Organizations change themselves by deciding about the premises of future decisions,

[27] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 69; Beer, *Heart of Enterprise*, 260 and, for the variety-matching statement, 235.

[28] Beer, *Diagnosing the System*, 86.

[29] Helmut Willke, *Systemtheorie entwickelter Gesellschaften: Dynamik und Riskanz moderner gesellschaftlicher Selbstorganisation* (Juventa, 1989); Helmut Willke, *Systemtheorie III: Steuerungstheorie; Grundzüge einer Theorie der Steuerung komplexer Sozialsysteme* (Gustav Fischer, 1995). *Kontextsteuerung* is Willke's term. The identification of it with Beer's principle of autonomy is ours: we have not found a place where Willke discusses Beer, and he should not be read as endorsing the equation.

which is not far from what we call design. And for Luhmann organizations are the only social systems that can act as units and be addressed as actors. That last point helps us.

The third is that the integration has already been done, and we mostly have not read it. Achterbergh and Vriens devote parts of one book to Luhmann and to Beer and argue that the two can be brought into a single framework. Johnson and Leydesdorff staged the confrontation directly, reporting Luhmann's position that businesses "do not 'live' as viable systems but operate because they relieve human beings from environmental complexity". Schwaninger and Groesser used operational closure and self-reference as *cybernetic* tools to explain organizational change.[30] Luhmann's own debts are cybernetic in any case: Ashby, von Foerster, Maturana and Varela, Spencer-Brown.[31] Reduction of complexity is a variety argument. The dispute is over the unit of analysis, not over whether regulation occurs.

The fourth is that the argument proves too much. If it holds, medicine, engineering and Luhmann's own sociology are equally self-description. A book about practice may fairly reply that a theory under which all intervention is impossible has not explained why some interventions reliably work.

What the objection does establish is a matter of grammar, and it is not trivial. It lands cleanly against an ontology, the claim that an organization *is* a viable system with five subsystems. It glances off a design language that offers a variety-based test of whether a set of governance arrangements can absorb the disturbances it faces. The cost to us is to stop making the first kind of claim, and to read the integration work before anyone in this tradition asserts incompatibility again.

Assemblage and network: the levels are not there

The next two positions attack the same structural feature from opposite directions, and both convert into amendments we should make.

Manuel DeLanda's assemblage theory insists on relations of *exteriority*: "a component part of an assemblage may be detached from it and plugged into a different assemblage in which its interactions are different", and

[30] Jan Achterbergh and Dirk Vriens, *Organizations: Social Systems Conducting Experiments* (Springer, 2009); Mark William Johnson and Loet Leydesdorff, "Beer's Viable System Model and Luhmann's Communication Theory: 'Organizations' from the Perspective of Meta-Games," *Systems Research and Behavioral Science* 32, no. 3 (2015): 266–82; Markus Schwaninger and Stefan N. Groesser, "Operational Closure and Self-Reference: On the Logic of Organizational Change," *Systems Research and Behavioral Science* 29, no. 4 (2012): 342–67. The quoted words about businesses that "do not 'live' as viable systems but operate because they relieve human beings from environmental complexity" are Johnson and Leydesdorff's own summary of Luhmann's position and not a quotation from Luhmann.

[31] W. Ross Ashby, *An Introduction to Cybernetics* (Chapman & Hall, 1956); Heinz von Foerster, ed., *Cybernetics of Cybernetics; or, The Control of Control and the Communication of Communication* (Biological Computer Laboratory, University of Illinois, 1974); Humberto R. Maturana and Francisco J. Varela, *Autopoiesis and Cognition: The Realization of the Living* (D. Reidel, 1980); George Spencer-Brown, *Laws of Form* (George Allen and Unwin, 1969).

"the properties of the component parts can never explain the relations which constitute a whole." A component's role in a larger whole does not define it.[32] System One, as we teach it, denies exactly this. A unit is an operational element because it produces what the organization exists to produce, and its function in the whole is taken to confer its identity. Those are relations of interiority, the organicist model assemblage theory was built to destroy. Beer's own text is less organicist than our teaching. He calls a subsidiary "a viable entity in itself" that belongs to its parent only "in a specially defined way", and he reports a divisional chairman "secretly moving to sell his division."[33] The model assigns a unit its role by function; its identity, on Beer's page, is its own, and DeLanda's point lands on how the role is assigned. The emblematic bad object of that tradition is the arborescent diagram, a root and branches, and the VSM diagram is usually drawn as one, although Beer wrote that its vertical dimension "has logical but not organizational implications."[34,35]

But DeLanda does not abolish levels. He builds an explicitly multi-scale social ontology in which persons, networks, organizations, cities and states are each real and irreducible. What he denies is that levels are internally *constitutive*. That is close to Beer's principle of autonomy, on which each System One is a viable system in its own right, with its own identity and purpose, whose cohesion with the whole is bought rather than assumed.[36] DeLanda's territorialization axis measures precisely how detachable a component is: cheap for a contractor, expensive for a licensed plant. The model does not assume uniform embeddedness in Beer's text: systems of concentrated purpose are "closely-knit," those of general purpose "more loosely coherent," and freedom is "in principle a computable function" of purpose.[37] Nobody has computed it, and our diagnoses treat embeddedness as uniform. The repair is still an index, now as the measure of a quantity Beer named and left uncomputed, and the model stays. A diagnosis should record how detachable each unit is, and stop treating a franchise, a subsidiary and a regulated utility as the same kind of part.

---

[32] Manuel DeLanda, *A New Philosophy of Society: Assemblage Theory and Social Complexity* (Continuum, 2006), 10–11. Both quoted sentences come from his statement of relations of exteriority in the opening chapter.

[33] Beer, *Diagnosing the System*, 1; Beer, *Heart of Enterprise*, 351.

[34] Gilles Deleuze and Félix Guattari, *A Thousand Plateaus: Capitalism and Schizophrenia*, trans. Brian Massumi (University of Minnesota Press, 1987), 3–25, originally published as *Mille plateaux* (Éditions de Minuit, 1980). The arborescent-rhizome opposition is theirs, not DeLanda's; he inherits it. We name them here because the text uses the term without naming them.

[35] Beer, *Heart of Enterprise*, 200. Against this, *Brain of the Firm*, 228, calls the same construction "the five-tier hierarchic model."

[36] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 158, for what Beer calls "the cybernetic argument for autonomy (or freedom)", which is what the text means by the principle of autonomy — Beer does not count it among his Principles of Organization — and 118 for the recursive system theorem; DeLanda, *A New Philosophy of Society*, for territorialization and deterritorialization as an axis of variation.

[37] Beer, *Heart of Enterprise*, 158.

Actor-network theory goes further and flattens everything. "Macro no longer describes a wider or a larger site in which the micro would be embedded like some Russian Matryoshka doll, but another equally local, equally micro place, which is connected to many others through some medium transporting specific types of traces"; "size and zoom should not be confused with connectedness."[38] Without a scale ordering, recursion↗ would have nothing to recurse over, if recursion were a ladder of sizes. Beer's is ordered by the containment of viable systems, not by size or zoom, and he warned against reading it as organizational level.[39] Latour's scale point therefore misses Beer's definition. What it hits is the claim that every level is a viable system in its own right.

Except that the oligopticon concedes Beer's mechanism. Latour's centers of calculation accumulate inscriptions and so act at a distance while seeing very little, very well. Cyberstride and the operations room are centers of calculation in precisely his sense.[40] Actor-network theory does not deny that they work; it denies that they work by being *above*. Beer's own account says a metasystem↗ works by attenuating, by throwing almost everything away. That is Latour's oligopticon in Ashbyan clothing, and the objection converts into a friendly amendment that takes Beer at his word. He called the recursive description "a vast attenuator of the variety proliferated by any large organization."[41] Keep recursion as that bookkeeping of who attenuates what for whom, and drop the ontological ladder our teaching has built on it.

Jackson: a placement, not a verdict

Michael Jackson's system of systems methodologies, first published with Paul Keys in 1984 and developed since, sorts approaches by the complexity of the system and the relations among participants: unitary, pluralist, and in the later versions coercive. It places organizational cybernetics in the complex-unitary cell. The characterization that gets quoted is that the VSM is a "unitary, functionalist" approach, "unable to deal with individual, social and political issues": useful where interests genuinely converge, untrustworthy where they do not, and in coercive contexts liable to make efficiency into the efficiency of domination.[42]

[38] Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford University Press, 2005), 176 and 187. Both quoted sentences belong to the "First Move: Localizing the Global" section of Part II (173–90).

[39] Beer, *Diagnosing the System*, 2; Beer, *Heart of Enterprise*, 308. Beer's criterion is set out in Chapter 10.

[40] Bruno Latour, *Science in Action: How to Follow Scientists and Engineers through Society* (Harvard University Press, 1987), chap. 6, "Centres of Calculation," 215–57, for centers of calculation; Latour, *Reassembling the Social*, 181–83, for the oligopticon, which sees "much too little" but sees it well (181). The application of either to Cyberstride and the operations room is ours, not Latour's; he does not discuss Cybersyn.

[41] Beer, *Heart of Enterprise*, 308.

[42] Michael C. Jackson and Paul Keys, "Towards a System of Systems Methodologies," *Journal of the Operational Research Society* 35, no. 6 (1984): 473–86. The 1984 grid has two axes only:

Two things about this are usually got wrong, and both errors are ours. The first is the tone. Jackson's 1988 paper on Beer is titled an *appreciation*, his 2019 treatment devotes a section to the value of organizational cybernetics to managers, and his stated project has always been complementarism rather than elimination. A chapter that reads him as an executioner will be corrected in review, and will deserve it. The second is that the exchange is live and recent. Angela Espinosa, writing in a Festschrift for Jackson, argues that he has become more supportive over the decades and that Beer's successors have addressed several of the original criticisms, and she proposes a posture she calls critical and empathetic. Other recent work has taken the power criticism seriously enough to graft an explicit power framework onto the model. Beer himself published his own account of the model's provenance and pathology in the same 1984 volume that carried Jackson and Keys.[43]

In one place the classification is weaker than it looks: "unitary" may be a category error. The principle of autonomy and the recursive system theorem make each System One a viable system with its own identity and purpose. The model *assumes* divergent sub-unit purposes; that is why the anti-oscillation function and the resource bargain exist at all. Whether a resource bargain amounts to pluralism is a real question and not yet a settled verdict. And a grid that sorts methodologies into cells is itself an exercise in boundary judgment, with premises of its own that Ulrich's

---

systems, simple or complex, and decision-makers, unitary or pluralist. Coercive contexts are named in that paper and deliberately left outside the grid (483, the page Jackson himself cites for it in 2019), on the ground that no systems methodology then addressed them, so the coercive column belongs to the later restatements — Michael C. Jackson, "Beyond a System of Systems Methodologies," *Journal of the Operational Research Society* 41, no. 8 (1990): 657–68; *Systems Thinking: Creative Holism for Managers* (John Wiley & Sons, 2003); and *Critical Systems Thinking and the Management of Complexity* (John Wiley & Sons, 2019), which further adds a "complicated" row and renames the horizontal axis from participants to stakeholders (163–64). We have not been able to fix from a primary source which publication carried the coercive column first: Jackson's 2019 account says only that "later renditions of the grid explicitly included 'coercive' problem contexts," and dates the grid's stable form to his *Systems Methodology for the Management Sciences* (Plenum, 1991) and *Systems Approaches to Management* (Kluwer/Plenum, 2000) (157–59). We have therefore left the text's "in the later versions" unqualified. The quoted characterization — a "unitary, functionalist" approach "unable to deal with individual, social and political issues" — is Angela Espinosa's rendering of Jackson's criticism, and is the form in which it circulates; see the next note.

[43] M. C. Jackson, "An Appreciation of Stafford Beer's 'Viable System' Viewpoint on Managerial Practice," *Journal of Management Studies* 25, no. 6 (1988): 557–73; Jackson, *Critical Systems Thinking and the Management of Complexity*, §13.6, "The Value of Organizational Cybernetics to Managers," 339–40; Angela Espinosa, "Revisiting the Viable System Model as an Emancipatory Systems Approach," *Systems Research and Behavioral Science* 42, no. 1 (2025): 171–88, in the festschrift introduced by Ashish Dwivedi and Gerald Midgley, "Festschrift for Mike Jackson," *Systems Research and Behavioral Science* 42, no. 1 (2025): 3–10; Nahed T. Zeini, "The Viable System Model Through a Power-Relations Lens: Towards a Power-Aware Conceptual Framework," *Systems Research and Behavioral Science* 43, no. 3 (2026): 1133–43; Stafford Beer, "The Viable System Model: Its Provenance, Development, Methodology and Pathology," *Journal of the Operational Research Society* 35, no. 1 (1984): 7–25. Espinosa's own phrase is a "critical empathetic approach." Beer's paper is in issue 1 and Jackson and Keys's in issue 6 of the same volume 35, which is what "the same 1984 volume" means; they are not in the same issue.

twelve questions would find. What we owe here is modest and overdue: state the placement, and stop presenting the model as a method to be executed in every context.

Which of these land

Three, and they land on the same place.

Ulrich lands, and the concession is the boundary argument Part One has already made: the first cut is somebody's, and we do not publish whose. Hayek's measure problem lands, and it is Ulrich's third question in an economist's register: the indices are chosen before the surprise, by the person who did not anticipate it. Cooper's point about the boundary lands, and it is the same concession a third time, stated as an act instead of a question or a measure. One defect, three descriptions: an act Beer named and never gave a procedure, and no reply from us to any of them.

The rest price amendments. Luhmann's objection lands against an ontology and glances off a design language; the cost is to stop making ontological claims and to read the integration work that already exists. Assemblage theory buys a territorialization index, the measure of an embeddedness Beer said varies and our diagnoses treat as uniform. Actor-network theory buys the demotion of recursion from a ladder of being to a record of who attenuates what for whom. Jackson's classification places the model and does not refute it, and its author has been saying so for decades. Process philosophy, minus Cooper's boundary point, describes a problem the Three-Four homeostat was built for.

The priced receipt

Receipt — publish the boundary judgments · Require every published diagnosis to answer Ulrich's twelve questions in both *is* and *ought* modes, including question 10, and to publish the boundary judgments beside the system map. Retrofitting one completed diagnosis is a five-day desk exercise; adding the protocol prospectively is roughly two interview-days per engagement. If competent analysts cannot answer the questions from the engagement record, the diagnosis is not reproducible enough to support its boundary claims. If the answers expose affected parties without a witness role, the design must record that as a defect rather than hide it in the system boundary.

That receipt addresses the external critics, and only the first of the three wounds. It buys nothing against the measure problem, which needs evidence rather than a protocol. Nor does it buy anything against Cooper, who is answered only by examining the first cut, not by documenting it after the fact. The hardest political charge is internal, what this tradition itself says about freedom and control, and Part Four holds it open rather than pretending this chapter has paid it.

Part III

False Friends and Enemies

The Corner

> Judy's corner is crowded. Between rounds the people in it lean over the ropes with advice, and she has begun to notice that almost none of it is about winning the fight in front of her. One tells her she has never been beaten, which is true and means nothing: nobody has ever scored her. One has a stopwatch and a clipboard and a method, and has been timing the wrong arm for years. One says that scoring is the wrong frame entirely, that what she does is a dance and a dance cannot be lost. Across the hall, in the seats, sit the people who say she is not a boxer at all. They have never watched a round. What frightens her is not them. It is that nobody has ever bought a ticket.

Parts One and Two brought charges. Part Four will say what survives them, and Part Five what it would cost to put the surviving claim at risk. First comes the question a reader of Chapter 9 is left with: if the defects never cost us anything, who was standing around not charging us?

The answer divides in two, and the two halves of this part follow the division.

The first half is about the communities that appear to be on our side. Some of them are. Checkland's recoverability and Ormerod's critical rationalism (Chapter 20), Troncale's call for validation (Chapter 15), and the free energy principle's public split between a principle and its process theories (Chapter 21) are all positions this book can adopt without amendment. The chapters here are about the others: neighbors whose support is real and whose intentions are not in question, but whose epistemic settlement, once we adopted it, made the model unchargeable. A community that tells us our model is useful rather than true has not attacked us. It has handed us the sentence we use to avoid the question, and we have used it. The first to hand it over was our own founder. "So remember: / a model is neither true nor false: / it is more or less useful," Beer wrote in 1985, in *Diagnosing the System for Organizations*.[44] In 1966 he had tied the same idea to a test: the usefulness of a model is measured "by straightforward testing of the predictive value of the scientific model. If this predictive value is low, we had better start again."[45] The tradition kept the sentence and dropped the test, and when the neighbors offered their own versions of the sentence, we took those without their tests as well. Beer also made the other kind of claim: "if a system has all of the necessary and sufficient conditions of viability, which we set out to discover and have now listed, it will live."[46]

[44] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2, where the three lines are set apart from the text, after "you are not determining absolute facts: you are establishing a set of conventions." Chapter 20 gives Beer's other versions of the sentence.

[45] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113. Chapter 21 quotes the passage in full.

[46] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 408, checked 29 September 2026 in the calibre PDF. Chapter 21 quotes the same passage.

The line this part draws between the two kinds of claim runs through the founder's own books before it runs between us and our neighbors.

These are the false friends of the title. The title accuses nobody of bad faith; it describes what an inherited criterion does when it is inherited or borrowed without the practice that earned it.

The second half is about the adversaries. There are fewer of them than the reception suggests, and they are weaker than the reception suggests, and that is the finding, not a relief. Chapter 16 already took the three objections that land, Werner Ulrich's, Friedrich Hayek's and Robert Cooper's, and this part does not revisit them. What remains is a rival ontology with no test in it; a set of sweeping dismissals of systems thinking in general that were never aimed at this model in particular; a body of formal criticism directed at a theorem we borrowed rather than at anything we built; and a silence. The silence is the serious one, and the last chapter is about it.

Four notes on reading the part.

It is about us, not about them. Every community described here is doing something defensible within its own program. System Dynamics built its validation philosophy and then built the test batteries to go with it (Chapter 20). The American Society for Cybernetics has never claimed to be an evidence-producing body and does not pretend otherwise (Chapter 19). Radical constructivism has been explicit for forty years about what it means by viability (Chapter 19). None of them owes us a failure condition. We took their conclusions and left their obligations, as we had already done with Beer's prediction test, and that charge is ours to answer.

The standard is the one Chapter 16 set. A position counts as an adversary here when it identifies something the model needs in order to do the job it claims, cannot supply from its own resources, and has not been given by any amendment since. A position counts as a false friend when it supplies a criterion that the model then uses in place of a test, and the criterion cannot in principle settle the claim it is used to settle. The two failures are different, and they are priced differently.

The instrumental/constitutive line does the work. Chapter 23 will fix the narrow claim on this distinction; we need the distinction here first, because it is what makes this part tractable. Usefulness, recoverability, fit to context and boundary critique are adequate criteria for instrumental claims, claims about whether a diagnosis helps. None of them can decide a constitutive claim: that there are five subsystems, that they are necessary, that they are sufficient, that System Five corresponds to anything in a cortex. Wherever this part says a criterion has been misapplied, this is the misapplication it means.

Not everything here has been verified. The quotations, attributions and dates in these six chapters were first drawn from two research notes prepared in September 2026. Some have since been checked against the original sources: the material from SCiO's own website on 14 September 2026, and on 23 and 25 September the Metaphorum page, the Schwaninger and Scheef sentence, the Hoverstadt podcast transcript, Espinosa's post and

the pages of Pfiffner's dissertation. That first-hand evidence is concentrated in Chapters 17 and 18, and the footnotes say where it appears. Beer's own sentences, added on 29 September 2026, do not come from those notes; they are quoted from his books and were checked against the printed pages. Each chapter flags the items still unverified, and they are not few. A book about the difference between a claim and a receipt cannot print a sourced claim on the strength of a summary of a source. Every quotation below must be read against the original before this part goes to press, and the chapters say so where it matters most.

Chapter 17 takes the community closest to us and the warrant it publishes: in fifty years nobody has found a case in which the model failed. Chapter 18 takes the two most recent quantitative tests of the model and shows that neither was built so that it could lose. Chapter 19 takes the host discipline's official self-description, cybernetics as a way of thinking rather than a body of facts, which is the opposite of Beer's "collection of principles that systems cannot violate,"[47] and asks what a way of thinking can be charged with. Chapter 20 takes usefulness instead of truth, the criterion that has done the most damage by being the most respectable, and begins with Beer's own version of it and the test he attached. Chapter 21 takes the last shelter, the claim that testing the model is a category error, sets the status Beer declared for his model in 1966 beside the one Karl Friston declared for the free energy principle, and asks what each community did with the exposure its founder accepted. Chapter 22 turns to the adversaries and finds that the one that costs us most is not a person.

[47] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 24: "What truly is given in nature is a collection of principles that systems cannot violate if they are to maintain their identity. These principles are the content of the science of cybernetics." In *Think Before You Think,* ed. David Whittaker (Wavestone Press, 2009), 24, he speaks of "fundamental facts" about how large probabilistic systems work "which seem to be invariant to the fabric of the system concerned." Both checked 29 September 2026 in the calibre PDFs.

Chapter 17

The Market Test

The claim

Metaphorum's website states that after fifty years no one has been able to find a case study in which the Viable System Model has failed.[1] Schwaninger and Scheef, from a different tradition and with a different purpose, record the same absence from the other end: to their knowledge, unsuccessful or failing applications of the model have not been reported in the literature.[2] Read together, the two sentences close a loop. The reading is ours; neither source makes the joint claim. Failures are not reported, so none is found, so the record is clean, so there is nothing to report.

The community nearest to us publishes a warrant, and the warrant cannot be false. Underneath it is a mechanism that needs neither a rhetorical trick nor anybody's dishonesty: a selection effect operating on a professional practice. It would produce the same result if every practitioner in the field were scrupulous. Most of them are.

What "failed" would have to mean

The sentence "no case study in which the VSM has failed" has no truth conditions, because this literature, as far as we have found, has never defined what it would mean for a model to fail in an engagement.[3] Take the candidates one at a time, with what each would require us to have recorded.

[1] Metaphorum, "Viable System Model," https://metaphorum.org/staffords-work/viable-system-model (the page title is "Viable System Model"), consulted 23 and 29 September 2026: "After 50 years, no-one has been able to find a case study in which the VSM has failed: in all its many and varied applications it has provided insightful diagnosis and has directed organisational restructuring to create effective organisations, based on individual autonomy and creativity." The page is live and editable; an archived copy should be deposited before print.

[2] Markus Schwaninger and Christine Scheef, "A Test of the Viable System Model: Theoretical Claim vs. Empirical Evidence," *Cybernetics and Systems* 47, no. 7 (2016): 544–69, https://doi.org/10.1080/01969722.2016.1209375: "To our knowledge, unsuccessful or failing applications of the VSM have not been reported in the literature." Checked verbatim on 23 September 2026 on p. 6 of the online-first version in the University of St. Gallen repository; the page in the printed issue is not yet confirmed. The authors do not characterize the absence as publication bias; earlier drafts of the companion working paper described it as a "publication-bias admission," and that phrasing is our interpretation and must be marked as one wherever it appears.

[3] Search on 29 September 2026: OpenAlex, title and abstract, "viable system model" AND (fail OR failure OR failed OR unsuccessful), ninety-five records. By title, none proposes a

A diagnosis could fail descriptively: the analyst maps five systems onto an organization, and the mapping turns out not to hold. A function said to be present is absent; a channel said to carry variety carries none. Establishing that requires an independent re-mapping by someone who did not produce the first. That is Chapter 5's problem about the ruler, and we have never run it.

A diagnosis could fail predictively: the model names a pathology, the pathology is left unrepaired, and the predicted consequence does not follow. Establishing that requires the prediction to have been written down in advance with a date on it, which our engagement reports do not do.

An intervention could fail: the recommended structural change is made and the organization does not improve, or declines. Establishing that requires a measure of improvement stated and fixed before the change. That returns us to Ulrich's third question in Chapter 16, and it is not standard practice either.

Or the engagement could fail: the client stops the work. This one we can observe. It is the only one we can observe, and it almost never reaches print. Beer printed a few, among them four proposals, in four countries, in which "the impending crisis exploded before work could begin."[4]

Four meanings: three that the field has never instrumented, and one that it almost never publishes. Against that background the claim that no failure has been found reports nothing about the model. It reports the design of the observation system.

The market test, stated plainly

The practitioner network does have an operative test, and its honesty should be acknowledged before it is criticized. Patrick Hoverstadt has put it directly: if it doesn't work, they don't ask you back.[5] That is a real feedback loop. It runs fast, it costs the practitioner money when it fires, and it is more than most consulting frameworks submit to. The warrant is older than the network: Beer's own confidence in the model came, "subjectively speaking," from "the speed with which the diagnosis can be made."[6]

---

definition of a failed VSM application; abstracts and full texts were not read, so the claim rests on a title screen.

[4]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 350: "I made speed the essence of my proposals in four other countries — but in each case the impending crisis exploded before work could begin." Beer does not name the countries and reads the four cases as validating "both the hypothesis about crisis and the strategy of swift action." At 323 he reports a descriptive setback in the Chilean work: in the model of the economy used until late 1972, "two levels of recursion were mixed together in that preliminary model." Both passages checked 29 September 2026 in the calibre PDF (printed pages).

[5]Patrick Hoverstadt, interviewed by Peter Hayward, "Patterns of Strategy," *FuturePod*, episode 199, 25 August 2024, https://www.futurepod.org/podcast/ep-199, transcript on the episode page (without timestamps): "Because if it doesn't work, they don't ask you back, and they don't tell other people about you." Checked on the episode page on 29 September 2026.

[6]Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 148: "Subjectively speaking, confidence in the VSM as applied to societary systems derives not so

It is also selected on the outcome variable, in the direction that hides failure. A client who does not ask you back does not co-author the case study. The practitioner who loses the account writes nothing; the practitioner who keeps it writes the chapter. Every mechanism that could turn a negative result into a published finding runs the other way. Confidentiality agreements cover the engagements that ended badly, reputational cost attaches to reporting them, and no venue has asked for them. The founder did. Beer proposed "a small team of organizational pathologists," organized "more like a lifeboat crew," to reach a dying organization, resuscitate it if possible and, failing that, perform "a post mortem … before rigor mortis had set in," because much had been learned from systems "viable no more."[7] Nothing in the practice described in this chapter answers to the proposal. The market test is a test of the practitioner, not of the model, and nobody records it even as a test of the practitioner.

The same structure appears in the network's other warrants. Angela Espinosa writes that the model "has been tested in several types of businesses, public and not-for-profit organisations and communities".[8] The word "tested" there means "applied": the cases are ones in which the model was used, reported by the people who used it. Her 2025 reply to Michael Jackson's charge that the model is unitary and functionalist answers largely by pointing to recent applications.[9] That answers a question about usefulness. Jackson's question was about frame.

---

much from the fact that the pathology of the viable system can be investigated with ease, as from the speed with which the diagnosis can be made." Checked 29 September 2026 in the calibre PDF (printed page 148).

[7] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150, at the end of the essay on the viable system model. Beer adds that he had done such work himself, "but only as the result of being fortuitously present at the deathbed." The register proposed in Chapter 30 is the written form of the same idea.

[8] Angela Espinosa, "Viable System Model: A Theory for Designing More Responsive Organisations," *i2Insights*, 24 January 2023, https://i2insights.org/2023/01/24/viable-system-model/; the quoted words were checked on the live post on 25 September 2026. The wording is not in Espinosa's 2023 book, *Sustainable Self-Governance in Businesses and Society: The Viable System Model in Action* (Routledge), searched 23 September 2026 in the calibre copy of the first edition. The same warrant appears in her 2023 book: "Nowadays, there are hundreds of examples of VSM applications available worldwide, which shows that we have achieved some success" (61; checked 29 September 2026 in the same copy).

[9] Angela Espinosa, "Revisiting the Viable System Model as an Emancipatory Systems Approach," *Systems Research and Behavioral Science* 42, no. 1 (2025): 171–88. **Unverified:** the characterization above rests on the abstract only. Read the paper before printing any account of how it answers Jackson, and correct this paragraph if the body of the article makes a frame-level argument that the abstract does not show. The words "unitary and functionalist" are Espinosa's summary of Jackson, not a quotation from him. Jackson's 1988 appreciation (M. C. Jackson, "An Appreciation of Stafford Beer's 'Viable System' Viewpoint on Managerial Practice," *Journal of Management Studies* 25, no. 6 (1988): 557–73) concludes that "in functionalist terms the VSM passes the test of managerial significance with flying colours" and that doubts appear "once we step outside the functionalist paradigm" (570–71), but does not use "unitary"; Jackson's later book places the model in "complex-unitary" contexts (M. C. Jackson, *Critical Systems Thinking and the Management of Complexity* (Wiley, 2019), 334 and fig. 13.7, 335: "I would argue that the VSM is most at home with 'complex-unitary' contexts"); Espinosa's 2023 book attributes the functionalist classification to Ulrich (1981) and Jackson (1991) (*Sustainable*

José Pérez Ríos's pathology taxonomy repays the most attention, because it is the field's best piece of derived diagnostic machinery and shows the problem in its cleanest form. As Chapter 9 says, the taxonomy is generated from the structure; it is not assembled from observed breakdowns.[10] We have not been able to find a validation of it against failure data: a set of organizations that failed, independently classified, with the taxonomy asked in advance which pathology should be present. Until that exists, the catalog is a vocabulary with a derivation. That is a real asset, and it is not evidence.

The selection rule, in writing

Everything above is an inference about how a record comes to be clean. The inference turns out to be unnecessary. SCiO (Systems and Complexity in Organisation, the UK professional body for systems practitioners and the network through which most VSM practice in Britain is taught, accredited and sold) publishes the rule at the head of its case-study library:

> The following are case studies of successful application of systems methodologies.[11]

That is the inclusion criterion, stated by the community itself, in its own voice, on a public page. A failed application is not suppressed. It is out of scope by construction, and the sentence saying so is the first thing a reader sees.

Of the twelve entries we were able to read in full, none reports a failure, a negative result, an abandoned engagement, or a limitation of the method

Self-Governance in Businesses and Society, 62). Checked 23 September 2026 against the calibre copies (Wiley online version of the 1988 article; first edition of the 2023 book).

[10] José Pérez Ríos, *Design and Diagnosis for Sustainable Organizations: The Viable System Method* (Berlin: Springer, 2012), ch. 3, "Pathologies of Organisations," 141–70: structural pathologies (142–46), functional pathologies by System and for the whole system (146–66), and pathologies of information systems and communication channels (166–69). Pérez Ríos states that he lists them using "a conceptual outline similar to the one in Chap. 2 concerning the diagnosis and design of organisations" (141), and presents them as the pathologies "which most commonly occur" (141), but the chapter reports no frequencies and no test of the taxonomy against cases. Checked 23 September 2026 against the calibre copy of the 2012 Springer edition. A search on 29 September 2026 (OpenAlex, "Pérez Ríos pathologies viable system", first twenty-five results by relevance) found no study that tests the taxonomy against failure data; two titles on organizational pathologies (2019 and 2025) were not read. See also Chapter 9.

[11] SCiO, "Case Studies," https://www.systemspractice.org/casestudies. **Verified 14 September 2026:** the sentence is the first line of the page and is quoted exactly. The page continues: "Some are formatted and written up as case studies. Others are articles or talks that cover case study material." Archive a copy before print, since the page is live and editable.

used.[12] The count is what the criterion predicts, and the criterion is not hidden.

The community's efficacy claims point to this same library as their evidence base. SCiO's own page for the model states that the Viable System Model "has proven to be a scalable, super-fast, flexible, rigorous, approach" and that it "has been used successfully" for improving productivity, governance and resource matching, across "organisations of all types and scales from small teams up to whole nation states."[13] Each of those sentences is true of the case library as constituted. None of them could be false of it. Anything stronger the site does not say, anywhere we could find: no claim that the model has never failed, no count of applications, no years-of-use figure. The strongest single claim on the public site turns out to be a percentage, and Chapter 18 follows it to its source.[14]

One detail completes the picture, and it is small enough to be missed. The only artifact in this community named "The VSM Test" is a thirty-question multiple-choice quiz with an eighty-five percent pass mark, testing a reader's recall of a twenty-eight-page guide.[15] It tests the practitioner. Nothing in the tradition's public apparatus tests the model, and the one thing named as though it did tests whether you have memorized it.

Why this community and not another

A reader could take this chapter as an attack on Metaphorum and SCiO, and that reading would miss the point twice over.

It would miss it first because these are the people who kept the work alive. Between Beer's death and the present, this network ran the conferences, taught the courses, maintained the archive and did the client work, and that is why there is a tradition to stress-test at all. A tradition that had been

[12]Twelve entries were read verbatim, so the count is stated for the twelve and not for the library. On 29 September 2026 the page listed forty-five entries, sixteen of them tagged as VSM case studies; many are available only to members or on request from the author, and their full texts were not read, so the number reporting a negative result is not known. The claim in the text is deliberately limited to what was read.

[13]SCiO, "Why use the Viable System Model (VSM)?", https://www.systemspractice.org/why-practice-systems-thinking/vsm. **Verified 14 September 2026** for the phrases quoted. The same page carries the claim that "A study has shown a 71% correlation between organizational viability and survivability of existential crises," uncited in place; Chapter 18 traces that figure to its source and to what the source actually says.

[14]**Correction to the research note.** The note "Community positions on falsification of the VSM and cybernetics" (11 September 2026) attributes to a SCiO course page the claim that the VSM "is the best indicator of any organisation's chances of surviving or avoiding crisis." That sentence was not found anywhere on the SCiO site in a search on 14 September 2026, including the courses catalogue and the VSM pages. It may be on a course detail page not read, it may have been edited away, or it may be a compressed recollection of the 71 per cent claim discussed in Chapter 18. Do not cite it until it is located.

[15]SCiO, "The VSM Test," https://www.systemspractice.org/thevsmtest. **Verified 14 September 2026** and re-read 29 September 2026: two free multiple-choice quizzes hosted on ClassMarker, the basic one thirty questions with a stated passing score of 85 per cent, both "based on the VSM Guide" (Michael Frahm, *The VSM Guide*, v1.3, n.d., 28 pp., downloadable from the same site). No scoring rationale, item development or validation is described.

abandoned would have no defects worth cataloging, because nobody would be teaching them.

It would miss it second because nobody in the network has defended the claim under examination as an argument in print. It is a line on a website. Nobody sat down and reasoned their way to an unfalsifiable warrant. The sentence is the natural summary of fifty years of practice in which no instrument for recording failure was ever built, and to the person writing it, it reads as confident rather than evasive. That is what makes it a false friend rather than an adversary. An adversary's argument can be answered. A warrant of this shape cannot be answered, only replaced, and replacing it is cheap.

One further asymmetry bears on the next chapter. The 2016 St. Gallen test set out to falsify the model, and it reported a rejected hypothesis. It was not the first attempt: its authors name two earlier St. Gallen dissertations, by Benjamin Frost (2005) and Cristina-Ioana Crisan Tran (2006), which we have not read.[16] The only response from this network that we have read treats the 2016 test as corroboration. Espinosa's 2023 book says the study "demonstrated that the claim that the VSM provides the necessary and sufficient precondition for the viability of any organisation is well justified," and that its authors "could not falsify that claim on empirical grounds, which corroborates the effectiveness of the model." The rejected hypothesis is not mentioned.[17] Other authors from the network cite the study, and we have not yet read how. In the one case we have read, the community used the test by citing its conclusion.

[16] Schwaninger and Scheef, "A Test of the Viable System Model," 7 (online-first version, as in the notes above): with regard to attempts at falsification, "only very few endeavors of that kind have been undertaken: two dissertations have engaged in bold ventures that pioneered tests of the model on a broader basis"; and 9: "The outstanding merits of Crisan Tran (2006) and Frost (2005) are in making the VSM accessible to quantitative analysis for the first time." The two dissertations, as confirmed in the catalog of the Deutsche Nationalbibliothek on 29 September 2026: Benjamin Frost, *Lebensfähigkeit von Communities of Practice im organisationalen Kontext* (PhD diss., University of St. Gallen, 2005; no. 3120 in Schwaninger and Scheef's reference list); Cristina-Ioana Crisan Tran, *Beers Viable-System-Model und die Lebensfähigkeit von Jungunternehmen: Eine empirische Studie* (PhD diss., University of St. Gallen, 2006; no. 3201; Schwaninger and Scheef give the subtitle as "Eine empirische Untersuchung"). We have not read either dissertation. What they found is known to us only from Schwaninger and Scheef's summary; see Chapter 18.

[17] Angela Espinosa, *Sustainable Self-Governance in Businesses and Society: The Viable System Model in Action* (Routledge, 2023), 272, section "VSM Research Landscape," checked 29 September 2026 in the calibre copy of the first edition, which cites Schwaninger and Scheef nowhere else. Citation search on 29 September 2026: OpenAlex, works citing doi:10.1080/01969722.2016.1209375 (58 works; Crossref counts 44). Inclusion rule: authors active in Metaphorum or SCiO. Candidates found and not yet read for how they use the study: David Lowe, Angela Espinosa and Mike Yearworth (2020); two works by Espinosa, Jon Walker and co-authors (2023); Espinosa and Andrea Martinez-Lozada (2025); Michael Frahm and Hamid Rahebi (2020). Schwaninger's own later papers restate the conclusion (Chapter 18) and are not counted as replies.

What this chapter does not establish

It does not establish that the model has failed anywhere. No such case is presented here, and this book has none.

It does not establish that the practitioners are wrong about their engagements. They are almost certainly right that most of their clients were helped, and nothing here contradicts that.

What it establishes is that the field's strongest public claim is a statement about its recording practices wearing the clothes of a statement about its model, and that the two can be separated at a price small enough to name.

Receipt — define failure, then open the register. Publish a four-line operational definition of a failed VSM application, covering the descriptive, predictive, interventional and engagement senses above, with a rule for each saying what evidence would establish it, and a preamble saying which notion of viability the definition uses: Beer's capacity to maintain a separate existence, or bare survival (Chapter 18's receipt faces the same choice). Then invite the practitioner network to enter its own terminated and unsuccessful engagements against that definition, anonymized, in the register proposed in Chapter 30, with a twelve-month amnesty window and no requirement to name clients. This is Beer's post-mortem proposal, carried out on paper. Price: one afternoon for the definition, circulated for comment at one Metaphorum conference; the register cost is already booked in Chapter 30 and this adds a single intake form to it. If the register fills, the "no failed case" claim retires and the field gains its first negative evidence base. If it stays empty after twelve months with the definition in place and the amnesty offered, that is itself a finding, and a much stronger one than the claim it replaces: the absence would then be an observed absence rather than an unobserved one. Either outcome is worth more than the present sentence, and the present sentence is the one thing the exercise cannot leave standing.

Handover

This community never framed a test. The next chapter deals with the one group that did.

Chapter 18

The Test That Could Not Lose

The claim

In 2016 Markus Schwaninger and Christine Scheef set out to falsify the Viable System Model. A year later Michael Dominik Pfiffner defended a dissertation at Utrecht that rated 135 Swiss organizations against it. These two studies put numbers on the model's central claim. They were not the first to try. Schwaninger and Scheef build on two earlier St. Gallen dissertations, by Benjamin Frost (2005) and Cristina-Ioana Crisan Tran (2006), which they credit with making the model accessible to quantitative analysis for the first time.[1] We have not read either dissertation, and this chapter examines only the two later studies. Between them, those two did more than anything else we have read from the tradition's fifty years. We owe their authors that debt before we state the charge.

The charge is the same for both, and it concerns design, not competence. Neither study was built so that a negative result could cost the model its central claim, and neither concealed the fact. In the first it shows in the sentence that explains away the one hypothesis that failed. In the second it shows in the construction of the dependent variable and in the specification of the statistical model, both set out plainly in the text, and neither of which the author appears to have seen as a problem.

[1]Markus Schwaninger and Christine Scheef, "A Test of the Viable System Model: Theoretical Claim vs. Empirical Evidence," *Cybernetics and Systems* 47, no. 7 (2016): 544–69, https://doi.org/10.1080/01969722.2016.1209375, 7 (page numbers of the online-first version, as in the note on their framing below): with regard to attempts at falsification, "only very few endeavors of that kind have been undertaken: two dissertations have engaged in bold ventures that pioneered tests of the model on a broader basis"; and 9: "The outstanding merits of Crisan Tran (2006) and Frost (2005) are in making the VSM accessible to quantitative analysis for the first time." The two dissertations, as confirmed in the catalog of the Deutsche Nationalbibliothek on 29 September 2026: Benjamin Frost, *Lebensfähigkeit von Communities of Practice im organisationalen Kontext* (PhD diss., University of St. Gallen, 2005; no. 3120 in Schwaninger and Scheef's reference list); Cristina-Ioana Crisan Tran, *Beers Viable-System-Model und die Lebensfähigkeit von Jungunternehmen: Eine empirische Studie* (PhD diss., University of St. Gallen, 2006; no. 3201; Schwaninger and Scheef give the subtitle as "Eine empirische Untersuchung"). We have not read either dissertation. Schwaninger and Scheef summarize them on p. 7: Crisan Tran surveyed startups, and "her study did not falsify the model," with results that "were not unambiguous," though "System 4 has a significant positive influence on the viability of a startup"; Frost studied communities of practice at Siemens AG and found "a strong correlation" between the values for the structural components and the long-term existence of these communities; "the empirical evidence provided by them, however, is limited, mainly due to scarcity of data." On p. 9 they add that Crisan Tran's survey attained 104 completed questionnaires. These are their accounts, not our reading.

An earlier draft of this chapter called the first study the only test ever run. That was wrong twice over, and instructively so. The second study is not in the journals this tradition reads; it turned up on a practitioner body's resource page, not in a database. The two earlier dissertations we learned of only from the St. Gallen paper's own review of earlier work. A field that cannot locate its own empirical work makes in its bibliography the error this book charges it with in its theory.

The St. Gallen test, 2016

Schwaninger and Scheef state the standard at the outset: the falsification criterion is widely considered stronger than the confirmation criterion, and they will test the theory of the Viable System Model according to the former.[2] Their design is a survey, 261 usable questionnaires from respondents who were mostly senior managers, testing whether the presence and quality of the model's components predict organizational viability.

Viability is not measured directly. Organizational culture stands in for it, and the same respondents rate both the components and the proxy.[3] Of the eight hypotheses, the component hypotheses 1, 2, 4 and 5 are supported, and so are 6 and 7, which take the components together. Hypothesis 3, concerning System Three, is rejected; its coefficient is negative. Hypothesis 8 is not supported. The conclusion: they have not succeeded in a falsification of Beer's allegation that the model specifies the necessary and sufficient preconditions for the viability of an organization.[4]

---

[2] Markus Schwaninger and Christine Scheef, "A Test of the Viable System Model: Theoretical Claim vs. Empirical Evidence," *Cybernetics and Systems* 47, no. 7 (2016): 544–69, https://doi.org/10.1080/01969722.2016.1209375, 6: "The falsification criterion is widely considered to be stronger than the confirmation criterion. Therefore, we will test the theory of the VSM according to the former." Page numbers in the notes to this section are those of the online-first version (published online 12 August 2016, pages 1–26), read on 29 September 2026 in the copy in the University of St. Gallen repository; the pages of the printed issue have not been checked.

[3] Schwaninger and Scheef, "A Test of the Viable System Model," 9–10 and 18: 349 complete questionnaires were collected at the 5th Business Congress of the Alpen-Adria University, Klagenfurt, and 261 were usable; the respondents were "in their majority … higher and top managers, including chief executive officers and owners of companies" (9); eleven items on organizational culture serve "as a proxy for the viability of an organization" (10). A supplementary analysis of three financial indicators was "inconclusive" and not pursued (18). No marker variable or other test for common-method variance is reported; the fourth limitation the authors list is that the study "has relied on one and the same informant, respectively, for both independent and dependent variables" (18). On common-method variance: Philip M. Podsakoff, Scott B. MacKenzie, Jeong-Yeon Lee, and Nathan P. Podsakoff, "Common Method Biases in Behavioral Research: A Critical Review of the Literature and Recommended Remedies," *Journal of Applied Psychology* 88, no. 5 (2003): 879–903, https://doi.org/10.1037/0021-9010.88.5.879.

[4] Schwaninger and Scheef, "A Test of the Viable System Model," 17: "We have not succeeded in a falsification of Beer's allegation that the model specifies the necessary and sufficient preconditions for the viability of an organization." The authors count "significant results for six hypotheses; one was neither supported nor rejected, and one was rejected" (17). Table 2 (14) gives the coefficient for Component 3 as −0.059 ($p < 0.10$) for the full sample of 261.

Three features of that design come before the substantive point. None of them is unusual in management research, which is part of why they passed.

The proxy carries the whole test. If viability is operationalized as organizational culture, what was tested is the relation between the model's components and a culture measure. The inference to viability is then only as strong as the argument that culture indexes viability, an argument the paper would have to make and a reader should look for. Whether the model's constitutive claim survives depends entirely on that bridge.

The common source is a known hazard. When one respondent supplies both the independent and the dependent measures, correlation between them is expected on methodological grounds alone, whether or not the underlying relation exists. The literature on common-method variance was already large in 2016. The authors report no marker variable or post-hoc test for it. They list the single informant among their limitations and recommend different informants for the two sets of variables in future work.

With self-report on both sides, the test asks managers whether their organization has the five functions and whether it has a good culture. Managers who believe in the one are likely to report the other. The test examines a belief structure at least as much as an organization.

The move that matters

These three features are not why the chapter exists. Design limitations are ordinary and correctable, and most of the management literature reports them honestly. The reason is what happened to Hypothesis 3.

System Three failed its test. The paper responds that it would be naïve and premature to conclude that this component is less important than postulated by theory,[5] and attributes the failure to changed organizational forms and to measurement.

Both explanations may well be correct. What matters is when they arrived: after the result, with no rule set beforehand saying what a rejected hypothesis would be allowed to cost. Without such a rule, "changed organizational forms" and "measurement" can absorb any negative finding whatever, about any component, in any future study. They are the general-purpose absorbers that Chapter 9 says our vocabulary supplies on demand, turning up in the one place in the literature where they should have been ruled out in advance.

[5] Schwaninger and Scheef, "A Test of the Viable System Model," 18, among the limitations: "It would be naïve and premature to conclude that this component is less important than postulated by theory." The explanation by alternative organizational forms is at 17, where the authors rule out "deficient measures for Component 3" as an explanation and assume instead that "the measurements were not sufficient to deal with new situations." The coefficient is negative (table 2, 14).

A falsification test without a declared revision rule is a confirmation exercise in statistical dress. The authors' integrity is not in question: they reported the rejection instead of burying it, which is the harder thing. The design is in question, because it had no way of losing. Had all five component hypotheses been supported, the model would have been corroborated. One was rejected, and the model was corroborated. The conclusion did not depend on the data, and the paper's own framing sentence supplies the standard by which that can be said.

The same paper makes a second move that compounds the first. It sets aside the critiques by Ulrich and by Amey on the ground that neither is of the empirical type, being rather deductive and often normative rather than descriptive.[6] The Ulrich they set aside is *Critical Heuristics of Social Planning* (1983); Chapter 16 shows what his earlier critique of Cybersyn establishes, and it is more than a matter of taste. A study that admits only empirical testing as criticism, and then explains away the one negative result of its own empirical test, has left no channel open in either direction.

Pfiffner's test, 2017

The other study is a doctoral dissertation defended at Utrecht University on 20 November 2017: Michael Dominik Pfiffner, *System viability of organizations and the aetiology of organizational crisis: A Quantitative Assessment of Stafford Beer's Viable System Model*, 313 pages.[7] It is more ambitious than the

[6] Schwaninger and Scheef, "A Test of the Viable System Model," 6: "Neither type of criticism is of the empirical type, but rather deductive, and they are often normative, not descriptive." The Ulrich text cited is *Critical Heuristics of Social Planning* (1983), which the paper says "disparages the VSM as the epitome of a 'cybernetic logic' "; the Amey text is Lloyd R. Amey, *Corporate Planning: A Systems View* (Praeger, 1986), summarized as holding that the model is "too general." Neither book has been obtained; see the receipt in Chapter 9, which buys the reproducible canon count that would settle what the published critical literature actually contains.

[7] Michael Dominik Pfiffner, *System viability of organizations and the aetiology of organizational crisis: A Quantitative Assessment of Stafford Beer's Viable System Model* (PhD diss., Universiteit Utrecht, 2017), xxvi + 280 pp. (313 pages in the repository PDF, counting its cover sheet and the preliminary pages); defended 20 November 2017, promotoren S.G.L. Schruijer and J.P.P.E.F. Boselie; Utrecht University Repository, http://hdl.handle.net/1874/356772. **Read in full text 14 September 2026.** Every quotation in this section is from the dissertation and has been checked against it. Page numbers are to the dissertation's printed pagination, checked 23 September 2026 against the calibre copy of the Utrecht repository PDF (printed page = PDF page − 33 throughout the main text): situation vignettes, 144–45; requisite-variety items ("always sufficient to"), 165; V.1–V.3, 167; K.1–K.6 and the "lack of requisite variety" gloss, 168 (also 86); "do not add value to the analysis," 169; managers' denial, 171; "which did clearly (not) undergo," 177; group means differing at the 99 per cent level, 191; the threshold of 92, 192; "technically uncertain and materially undifferentiated," 197; the reflective specification, loadings 0.791–0.982, "necessary and sufficient invariants" and −0.701, 213; R^2 0.5043, 218; "anticipation of the dependent variable," 219. Data were collected between December 2014 and April 2015; the survey was conducted in German (179). The dissertation is listed on SCiO's page "Why use the Viable System Model (VSM)?" as the study behind its "71% correlation" sentence and is hosted as a free SCiO resource, https://www.systemspractice.org/resources/system-viability-organizations-and-aetiology-organizational-crisis-quantitative (both checked 29 September

2016 study and better designed exactly where that study is weakest, and it fails in a more interesting way.

What it does, and what it gets right

Pfiffner turns the viability conditions into a questionnaire with a method he developed for the purpose, the situation vignette: an idealized control situation, stated in plain business language, against which a real organization can be compared. The sample is 135 organizations in German-speaking Switzerland, split in advance into two groups, 74 that had undergone an existential crisis and 61 that had not. Crisis is defined narrowly, and the definition holds throughout: a slowly developing existential threat to the whole organization, with disasters, sabotage and force majeure excluded.

Three of his design choices come first, because they are better than this tradition's norm and better than the 2016 study.

The organizations are not rated by their own managers. Pfiffner rules that out explicitly: crisis is the worst thing that can happen to an organization, and denial is the natural response. The raters are 22 third parties with close professional knowledge of the organizations: 16 bankers, including recovery and special-financing officers, five consultants from the author's own firm, and one fiduciary. Each had to clear a threshold on a control variable for case knowledge.[8]

The raters did not know the theory. Only two of the 22 said they had known the Viable System Model before the survey. They answered questions about functional traits in business language, unaware that they were performing a cybernetic analysis. Whatever else the instrument measures, it does not measure the raters' belief in Beer.

And Pfiffner defines the word *predictive* instead of trading on it. In his hypothesis, he states, it "refers to the anticipation of the dependent variable out of independent variable scores": statistical prediction within one dataset, not measurement at one time and observation at a later one. An earlier draft of this chapter guessed that the word was carrying more

2026); that is the practitioner body's resource page the text refers to. Note that Michael Dominik Pfiffner is a different person from Martin Pfiffner, author of *Die dritte Dimension des Organisierens / The Neurology of Business*.

[8]Twenty-two respondents: 16 bankers identified through a search for recovery and special-financing officers and commercial bankers at Swiss banks, five consultants approached among current and former colleagues at the author's own employing firm, and one fiduciary; drawn from 13 banks, 2 consultancies and 1 fiduciary in German-speaking Switzerland. Case knowledge was controlled by a self-rating on the Swiss six-point school scale with a minimum of 5 required; the sample mean was 5.36 (175, 181). The 16/5/1 split and the author's firm are at p. 179; the sample description two pages later gives fifteen bankers and six consultants (181–82), an inconsistency in the dissertation itself. That only two raters knew the model: 176–77. Self-rating ruled out because of denial: 171. One point for the authors to weigh before citing this study in support of anything: five of the twenty-two raters came from the author's own firm, and the dissertation's exposition of the model leans throughout on Malik's representations. Neither fact is concealed and neither is disqualifying; both belong in any account that presents the raters as independent.

weight than the design could bear. We guessed wrong. The definition is in the text, where any reader who checks will find it.

The reported result is a path coefficient of −0.701 between the system viability construct and the crisis construct, with R^2 of 0.5043, and a bivariate finding that the two groups differ on the viability scores at the 99 per cent level. He also reports a threshold: at or below a viability score of 92 out of 176, every organization in the sample had experienced a crisis.

Where the test stops being able to lose

The trouble is in the dependent variable. It is set out plainly in the methods and again in the results, which is the only reason it can be reported here.

Crisis was meant to be a dummy: crisis yes, crisis no. Pfiffner writes that this "ultimately became technically uncertain and materially undifferentiated," and that a combination of crisis and inverted viability indicators was used instead. The final outcome construct is the inverse of three viability items together with three crisis items.[9]

The three viability items ask the rater to place the organization on a scale from "perishing" to "thriving"; whether it can adapt to rapidly changing circumstances; and whether it is "fundamentally in a position to maintain its existence for an unlimited period of time." These are judgments about viability. They sit on the outcome side of a regression whose predictor is viability, inverted. Whatever the −0.701 measures, part of it is the agreement of one rater with himself.

The items that left the model are the ones that could have made it fail. Pfiffner reports that "some indicators do not add value to the analysis, so ultimately a reduced set of indicators was used." The dropped items asked whether the organization ceased to exist as an independent unit, whether it had been unable to fulfil its core mandate, and whether it had been unable to meet its financial obligations. Demise, mission failure, insolvency: these are the hard ones, matters of record and not of judgment. What remained on the outcome side were judgments, made by the person who supplied the predictor, in the same sitting.

One pairing is close to definitional. On the predictor side, the requisite-variety items ask whether the organization's ability to react, scope of action and range of offerings "are always sufficient to … master expected" and then "even unexpected" issues or challenges. On the outcome side, the

[9]The questionnaire is reproduced in the dissertation's appendix. The final outcome construct combines the inverse of viability items V.1 to V.3 with crisis items K.1, K.2 and K.6; items K.3 (whether the organization ceased to exist as an independent unit), K.4 (inability to fulfil its core mandate) and K.5 (inability to meet financial obligations) were among those dropped. The reported basis for dropping them is that "some indicators do not add value to the analysis" (169; items at 167–68). Pfiffner also reports that the respondents' own initial crisis/non-crisis classification "was only a first indicator and was not used later in the analysis" (178), the outcome score being computed from the questionnaire items instead — which is what puts the inverted viability items into the dependent variable.

surviving crisis item K.6 asks whether "the demands faced by the organization massively exceed(ed) its capabilities, resulting in it spinning out of control." Pfiffner's own commentary glosses that item as "basically a lack of requisite variety." A predictor and its near-negation stand on opposite sides of the equation, and their correlation is reported as a finding about Beer's model.

Two further features limit what the number can carry. The raters were asked to pick from their own portfolios organizations "which did clearly undergo" or "which did clearly not undergo" an existential organizational crisis. That is extreme-group sampling, and it inflates any association it finds.[10] Interrater reliability rests on a single case: one organization was rated by two people, and their scores correlated at 0.73. Pfiffner reports this openly, and it is one doublet, not a reliability study.

The specification that made necessity unfalsifiable

The deepest problem lies furthest from the surface, and it decides what the study could have found.

The structural model is reflective and specified downward. Pfiffner states that "system viability causes the availability and VSM conformity of operations, of a metasystem and of the relevant VSM principles," and that these in turn cause the availability of Systems One to Five. That reverses the claim under test. The model asserts that the five functions are necessary *for* viability; the statistical model treats viability as a single latent factor of which the functions are interchangeable symptoms.

In a reflective model the indicators are expected to move together, because the specification assumes they will. The loadings duly come out between 0.791 and 0.982. Pfiffner reads them as follows: the path coefficients "stand for the necessary and sufficient invariants of organizational viability as defined by the VSM... they are of equal or largely similar importance. This corresponds to the theoretical claim of the VSM that if only one VSM characteristic is lacking, the system loses its viability."

The inference does not hold. Uniformly high loadings in a reflective model show the design working as specified; they are no evidence that each function is necessary. A test of necessity would need a formative measurement model, with the functions constituting viability instead of expressing it, and would then ask whether any one block carries weight the others cannot supply. As specified, no possible pattern of results could have shown a subsystem to be dispensable.[11] The claim this book calls

[10] Kristopher J. Preacher, Derek D. Rucker, Robert C. MacCallum, and W. Alan Nicewander, "Use of the Extreme Groups Approach: A Critical Reexamination and New Recommendations," *Psychological Methods* 10, no. 2 (2005): 178–92, https://doi.org/10.1037/1082-989X.10.2.178.

[11] The reflective, downward specification is stated in the results chapter and the interpretation of the loadings is quoted in the text (213). Variance inflation factors are reported as exceeding a conservative threshold of 5 while remaining below 10 (211–12; the highest is 7.65) — consistent with the near-uniform loadings a single-rater reflective design would produce. **For Part Five:** the dataset could in principle answer the necessity question under a

constitutive was placed, in the one study that took it as its subject, where the statistics could not reach it.

A formative specification would still leave one reply open, and the reply comes from Beer. In his texts a function is not simply missing. System Two "is always present, or the organization would shake itself to pieces"; "System Four exists — it has to: but it is disseminated, and lacks integration"; and an enterprise whose metasystem has collapsed into System Three is a "decerebrate cat" that "lives on; its viability is ensured by a bogus environment."[12] Pfiffner's sentence that "if only one VSM characteristic is lacking, the system loses its viability" is his gloss. Beer's texts treat a deficient function as a pathology, not as the end of the organization. So a study that finds a surviving organization without a recognizable System Four can be answered from the founder's page: the function is there, unofficially. A test of necessity has to close that reply before the data arrive. It needs each function coded as official, disseminated or absent by indicators registered in advance, and an outcome that indexes viability in Beer's sense, the capacity to maintain a separate existence, and not bare survival. The coding in Chapter 28 sets this out.

The loop, inside the study that adopts Popper

Pfiffner states his epistemology at the outset. He follows Popper, he writes, in holding that all knowledge is hypothetical and "cannot be verified but only falsified," and he names critical rationalism as his basis.

In the discussion chapter, among the merits of the model he is using, he lists this: the VSM "has a long track record of application and (peer) review and, to date, has not been falsified." The claim appears earlier too, quoted from Schwaninger as part of the warrant for the project: the model "has not been falsified until today".[13]

---

formative re-specification, which would be the cheapest decisive test this book has identified. Chapter 29 sets out that re-analysis, and Chapter 30 prices it. On the difference between reflective and formative measurement models: Cheryl Burke Jarvis, Scott B. MacKenzie, and Philip M. Podsakoff, "A Critical Review of Construct Indicators and Measurement Model Misspecification in Marketing and Consumer Research," *Journal of Consumer Research* 30, no. 2 (2003): 199–218, https://doi.org/10.1086/376806; Scott B. MacKenzie, Philip M. Podsakoff, and Cheryl Burke Jarvis, "The Problem of Measurement Model Misspecification in Behavioral and Organizational Research and Some Recommended Solutions," *Journal of Applied Psychology* 90, no. 4 (2005): 710–30, https://doi.org/10.1037/0021-9010.90.4.710.

[12] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 66; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 143 and 265. At 265 the decerebrate cat "lives on; its viability is ensured by a bogus environment, and is sustained by artificial sustenance. If you prod its leg, it kicks back. And this is called 'living'." Beer's scare quotes around "living" mark the state as persistence rather than health, although the same sentence speaks of "its viability." Pfiffner's gloss is quoted from the dissertation at 213.

[13] Pfiffner states his critical-rationalist position in the introductory chapter (8), following Popper and Keuth, *Logik der Forschung*, 11th ed. (2005). The "has not been falsified" claims appear in the theoretical framework, quoting Schwaninger (2000) (85), and again in the discussion of findings in the author's own voice (222). **Verified 14 September 2026.**

So Chapter 9's charge turns up where it should be least at home. A researcher who adopts falsificationism in his Chapter 1 cites, in his Chapter 6, the absence of a falsification as a point in the model's favor. Yet the absence is what his own discipline's failure to look had produced, and his own study was built so that the model's central claim was not among the things that could fail. Nothing here is dishonest. The loop simply lies in the literature, ready for anyone who wants a reason to trust the model, and it closes without anyone noticing.

What Pfiffner establishes, and what he does not

He establishes something real, and we do not want to take it away. Judgments about an organization's control functions, made in plain language by informed outsiders who do not know the theory, separate organizations that went through an existential crisis from organizations that did not. That is a finding, the first of its kind among the studies we have read in this tradition, and more than the field had before.

Beer hoped for a result of this kind. Of his own identity exercise, in which managers answer fifteen questions built on the model, he wrote that "maybe (one day) thresholds will be established, to exceed which will predict collapse."[14] Pfiffner's score of 92 is a threshold of that kind, found within one dataset. Beer's verb was *predict*, and Pfiffner, by his own definition of the word, did not claim it: a threshold that predicts collapse needs the measurement to come before the collapse.

What it does not establish is the model's constitutive claim. It does not show that five functions are necessary, or that they are sufficient, or that any one of them carries weight another does not. It does not distinguish the Viable System Model from any other instrument that asks informed observers whether an organization seems well run, because no rival instrument was tested against it. And an outcome variable that contains the predictor inflates the size of the association.[15]

[14] Beer, *Heart of Enterprise*, 432, on the cluster analysis of managers' answers to the identity exercise (431–33). At 433 Beer reports the first trial: twenty managers who "cheerfully asserted that they shared a corporate identity" were separated by the clusters into their two divisions, an early instance of self-report checked against its own internal structure.

[15] The dissertation's own limitations section (244–46) names three: the abstractness of VSM theory, the level of detail the questionnaire could carry, and sampling — which it describes as "something between random and convenience sampling" (245), concluding that the results "are not generalizable in the strongest sense of statistical requirements" (246). It does not raise the construct overlap in the dependent variable, the single-rater design, the extreme-group selection, or the consequences of the reflective specification. The four candid limitations quoted in earlier drafts of this chapter came from the author's talk "Predicting Organizational Crisis using the VSM," SCiO Open Meeting – Autumn 2018, Manchester, 22 October 2018 (slides and video on SCiO's resource page, https://www.systemspractice.org/resources/predicting-organizational-crisis-using-vsm, marked "Contact the Author"), not from the dissertation; where the two differ, cite the dissertation.

What the field did with it

The reception tells as much as the study. The one response from the practitioner network we have read cites the study as corroboration (Chapter 17). Schwaninger restated the conclusion in 2019.[16] Michael Jackson's 2019 textbook reported that the results "largely supported the hypotheses" and added that the authors "are, nevertheless, keen to point out that all they can definitively claim is not to have falsified the VSM."[17] In 2025 a PRISMA-protocol systematic review of the VSM literature reported limited empirical validation and a lack of empirical studies, with most research remaining conceptual or simulation-based. It does not cite Schwaninger and Scheef at all.[18]

That last point survives a reading of the review's reference list, and it is the sharpest single fact in this part. A systematic review, run under a protocol whose purpose is exhaustive retrieval, surfaced neither the 2016 test nor Pfiffner's dissertation. Either the search missed them, which says something about how the literature is indexed and named, or they were screened out, which says something about what counts as a study of the VSM. Both readings should trouble us, and neither can be settled without opening the review's search log.

What this chapter does not establish

It does not establish that System Three is unimportant, or that the 2016 result tells against the model. A single rejected hypothesis in a survey with a proxy dependent variable establishes very little in either direction. We concede that freely; it is the authors' own methodological point, turned around.

Nor does it establish that a better-designed study would falsify the model. Such a study might corroborate it more convincingly than this one did, and that outcome would be worth having.

---

[16] Markus Schwaninger, "Governance for Intelligent Organizations: A Cybernetic Contribution," *Kybernetes* 48, no. 1 (2019): 35–57, https://doi.org/10.1108/K-01-2018-0019. In the accepted manuscript in the University of St. Gallen repository (read 29 September 2026; journal pages not checked) he summarizes the 2016 study and concludes: "we did not succeed in falsifying the VSM. At the contrary, the evidence confirms the model's theoretical claim as justified undoubtedly."

[17] Michael C. Jackson, *Critical Systems Thinking and the Management of Complexity* (Wiley, 2019), 327, checked 29 September 2026 in the calibre PDF (printed page). Jackson's account does not mention the rejected hypothesis on System Three.

[18] Mohammad Reza Fathi, Soraya Birami, Alireza Payvar, and SeyedAli Doorafshan, "A Systematic Review of the Viable System Model: Applications, Insights, and Future Directions," *Journal of Systems Thinking in Practice* 4, no. 3 (2025): 109–45, https://doi.org/10.22067/jstinp.2025.91552.1136: "limited empirical validation" (abstract, 109) and "the lack of empirical studies, as most research remains conceptual or simulation-based" (139–40). Its reference list of thirty-eight items, read on 29 September 2026 in the journal's PDF, cites neither Schwaninger and Scheef (2016) nor Pfiffner (2017); its only work by Schwaninger is "Design for Viable Organizations" (2006).

It does not establish that Pfiffner's result is wrong. The result may be right. It is an association within one dataset, not a prediction over time, but if it holds, this tradition has spent nine years not noticing what may be its best evidence, which is an indictment of its own kind.

What it does establish is that both attempts to test the founder's stated claim, that the model gives the necessary and sufficient conditions of viability, went ahead without the component that makes a test bite, and that the missing component is a sentence written in advance.

Receipt — the pre-registered replication, and a rule stated first. Replicate the 2016 design with four changes: a viability measure independent of the respondents' own reports, drawn from records rather than opinion and declared in advance as an index of viability in Beer's sense rather than of bare survival (survival, solvency, or continuation of the organizational unit over a defined window are the record-based floor of such a measure, not the whole of it, since Beer himself holds that an organization which is there is not thereby "effectively there"[19]); separate sources for the component ratings and the outcome; each function coded official, disseminated or absent by indicators registered before fieldwork, under the coding in Chapter 28; and, before any data are seen, a registered revision rule naming which components, at what effect size and with what confidence, would be demoted from necessary to contributory. Price: one registered report at a journal with a results-blind first stage — neither *Systems Research and Behavioral Science* nor *Kybernetes* offers the format at present, so the track would have to be proposed to one of them or placed with a journal that does[20] — plus one doctoral project's fieldwork, or a secondary analysis of the 2016 dataset if the authors will release it, which would cost a fraction and could be requested this year. If the replication supports the components under a rule that could have demoted them, the model gains its first corroboration that means something. If a component fails under the declared rule, the rule is honoured and the claim is narrowed in print. The test is worth running only because both outcomes are payable, and neither is payable today.

Handover

Chapters 17 and 18 took the two communities that own the model. The next three take the neighbours whose criteria we borrowed, beginning with the host discipline itself.

[19] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 148: "The fact that it is there does not prove that it is effectively there, witness universities, nor efficiently there, witness hospitals."

[20] Center for Open Science, list of journals that have adopted Registered Reports, https://www.cos.io/initiatives/registered-reports, checked 29 September 2026: neither journal, nor *Cybernetics and Systems*, is listed. On the format: Christopher D. Chambers, "Registered Reports: A New Publishing Initiative at Cortex," *Cortex* 49, no. 3 (2013): 609–10, https://doi.org/10.1016/j.cortex.2012.12.016.

Chapter 19

A Way of Thinking Has No Failure Condition

The field's official self-description

The ASC Declaration of September 1983, drafted by Ernst von Glasersfeld and prefaced by Stuart Umpleby as past president, is still published as the society's statement of what cybernetics is. Its first sentence reads: "Cybernetics is a way of thinking, not a collection of facts." Further on it quotes von Foerster: "Objectivity is a subject's delusion that observing can be done without him."[1]

One caveat decides how this chapter can use that document. The American Society for Cybernetics is not the Viable System Model community. Using its documents to indict Metaphorum would be a category slip, and a reviewer would say so on first reading. The claim here is the narrower one: these documents show what the *host discipline* rewards, publishes and archives. The VSM community's own norms are the subject of Chapter 17, and they are argued there from that community's own material.

The narrower claim matters because the host discipline's self-description does more than set the background. It is the sentence the tradition reaches for whenever someone raises the question of testing, and it has an institutional home, a named author and a date.

Take the declaration at face value, as it is meant. A way of thinking has no failure condition. No observation could meet it and defeat it, because it makes no claim an observation could contradict. The declaration is not evasive and defends nothing; it gives an accurate description of one of the two things cybernetics is. The difficulty is that the field is also the other thing. The declaration knows this, since it describes "two major orientations" that "have lived side by side in cybernetics from the beginning," one towards the design of technology and one towards knowledge, but it does not say which of its claims belong to which. Beer had put the other thing in

[1] American Society for Cybernetics, "Declaration of The American Society for Cybernetics," September 1983, written by Ernst von Glasersfeld, with a preface by Stuart A. Umpleby, "ASC President 1980-1982," https://asc-cybernetics.org/the-field-of-cybernetics/, accessed 29 September 2026. The von Foerster passage is in the Declaration's paragraph on the philosophy of science, not at its close. The same text states that "knowledge must not be taken to be a picture of objective reality but rather as a particular way of organizing experience," and it describes "two major orientations" in cybernetics, one technological and one concerned with knowledge.

a sentence of his own: what is given in nature is "a collection of principles that systems cannot violate if they are to maintain their identity," and these principles "are the content of the science of cybernetics."[2]

The society's Definitions page sets out about forty definitions side by side, among them Wiener's, McCulloch's, Beer's, Bateson's and Pask's, and declines to adjudicate between them. Larry Richards compiled the list in 1987, when he was the society's president, and in 1999 stated its intent: "to demonstrate that one of the distinguishing features of cybernetics might be that it could legitimately have multiple definitions without contradicting itself," and to stimulate dialogue on the motivations of those who proposed them.[3] A named officer of the society thus sets out, as a goal, that the field need not resolve its own terms.

On the same page sits Beer's own joke about the impossibility of defining cybernetics: the condemned professor whose last request is to answer the question definitively, and the doctoral student who asks to be executed second.[4] It is a good joke. Told by Beer in 2001, it is worth more to this book than any outside critic making the same point in earnest, and it should be read the way he told it, as coming from a man who knew exactly what he had and had not pinned down.

What the field rewards

Institutions show their standards in what they honor. The ASC's stated award criteria are prominence in transdisciplinary research, impact of applications, service to the society and community, excellence in education and pedagogy, and advocacy for the field.[5]

---

[2]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 24. In *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 24, he speaks of "fundamental facts" about how large probabilistic systems work "which seem to be invariant to the fabric of the system concerned." Both checked 29 September 2026 in the calibre PDFs. The Declaration's "two major orientations" paragraph was read on the live page the same day.

[3]Larry Richards, "1987 ASC Compilation of Cybernetics Definitions," dated October 1999, on the ASC "Definitions" page, https://asc-cybernetics.org/definitions/, accessed 29 September 2026: "In 1987, when I was ASC president, I developed a list of definitions/descriptions." The second aim reads "to stimulate dialogue on what the motivations (intentions, desires, etc.) of those who have proposed different definitions might be."

[4]ASC "Definitions" page, "A Joke Related by Stafford Beer, October 2001," linked there to its source, Beer's address "What Is Cybernetics?" on receiving an honorary doctorate from the University of Valladolid (typescript formerly at nickgreen.pwp.blueyonder.co.uk, read 29 September 2026 in the Internet Archive copy of 4 March 2016; the joke is on its second page). The address was published as Stafford Beer, "What Is Cybernetics?," *Kybernetes* 31, no. 2 (2002): 209–19; the *Kybernetes* pagination of the joke has not been checked. The same page also carries a line attributed to Shannon advising Wiener to use the word precisely because nobody knows what it means; the site itself labels that as widely quoted and merely attributed. Treat it as an anecdote and do not build anything on it.

[5]ASC, "ASC Awards," https://asc-cybernetics.org/awards/, accessed 29 September 2026. The page lists the five kinds of contribution in these words: "Prominence in transdisciplinary research for cybernetics and systems," "Impact of applications of cybernetic theory and practice," "Service to the ASC and the cybernetics professional community more broadly," "Excellence in education and pedagogy for cybernetics and systems theory and practice," and

Empirical test, replication, prediction, having been shown wrong and having revised: none of them is among the criteria. That is a cheap and checkable receipt for the incentive argument in Chapter 9. A field that honors advocacy and application impact, and does not honor refutation, will not produce refutations. Nobody designed it to have that effect. It has it anyway.

The society's current activity points the same way. As of 29 September 2026 its initiatives page lists a speaker series, the #NewMacyMeetings, a student-run study group and two working groups, one on the society's archives and one on art, media and cybernetics; the 2026 conference theme was conversational confluences, and the society's own strapline describes cybernetics as an interdisciplinary conversation.[6] It has no working group on application, evidence, method or management. Sixty-two years after its founding, the organized capacity of American cybernetics goes into conversation, pedagogy and art. The observation carries a date because it will change, and it should be printed with its date and its URL.

Two absences on the site bear directly on the program in Part Five. The society publishes no methods guidance, no replication or data-sharing norms, and no register of applications and their outcomes; and it keeps no page collecting criticism of the field or of its models.[7] The live register proposed in Chapter 30 has no counterpart at the discipline's oldest American society.

The constructivist settlement, and one word

The second-order tradition is not hostile to testing; it relocates it. Umpleby holds that conjectures and refutations are the current consensus and that second-order science should add the observer, the effect of theories

"Advocacy for the cybernetics community and to advance and promote cybernetics." The awards were restructured in 2024. The Wiener Medal is "awarded for high achievement in transdisciplinary research in cybernetics and systems"; the McCulloch Award is for "significant impact in the community," through real-world applications or service to the society; ASC Fellows are recognized for "a lifetime career achievement." The published lists run back to 1968 (Wiener Medal: Warren McCulloch and Robert C. Wood); entries before 1983 are few.

[6]ASC, "Initiatives," https://asc-cybernetics.org/initiatives/, accessed 29 September 2026: Speakers Series, #NewMacyMeetings, Archives Working Group, Art, Media, and Cybernetics (AMC) Working Group, ASC Study Group ("Hosted by a student, for students"). The site's strapline reads "Advancing and promoting cybernetics as an interdisciplinary conversation." 2026 conference: *Conversational Confluences*, 3–7 August 2026, Ouro Preto, Brazil, https://events.asc-cybernetics.org/2026/. Founding: "The American Society for Cybernetics was founded in 1964," ASC, "ASC History," https://asc-cybernetics.org/asc-history/. Time-stamped: re-check immediately before print and update the date.

[7]Search path, 23 and 29 September 2026: the site's navigation from https://asc-cybernetics.org/ and the pages /the-field-of-cybernetics/, /definitions/, /awards/, /initiatives/, /asc-history/, /journals/, /books/, /affiliated-publications/ and /publications/ (the archive of newsletters and collections). None carries methods guidance, replication or data-sharing norms, a register of applications and outcomes, or a page collecting criticism. Pages not linked from these were not searched.

on what they describe, and methods as contributions.[8] Ranulph Glanville read conjecture and refutation as a circular design conversation, and he argued that it is inappropriate to require design to be "scientific," since scientific research is itself a restricted form of design.[9] Neither position forbids what this book asks for. Both make it somebody else's business.

The expected objection from this quarter is that recursion levels depend on the observer, and it is a good objection. Chapter 29 converts it into a test: analyst agreement on the same organization, with and without a declared purpose. An observer-dependence claim that can be operationalized as a measurable agreement rate is a hypothesis; one that cannot is a position. Which of the two ours is, we have never said.

Radical constructivism supplies the sharpest instance, and it is a word rather than an argument. For von Glasersfeld, knowledge is viable as long as it is useful in accomplishing a task or achieving a goal: viability is an epistemic property of a model, held by a knower.[10] Beer's viability belongs to the system. He takes the dictionary's "able to maintain a separate existence" and adds that an organization "is viable if it can survive in a particular sort of environment."[11] But on his own account an observer decides which collection counts as the system. "A system is not something given in nature, but something defined by intelligence"; two people "will

[8] Stuart Umpleby, "Second-Order Science: Logic, Strategies, Methods," *Constructivist Foundations* 10, no. 1 (2014): 16–23. Checked against the published abstract only (29 September 2026): "there is widespread agreement on how to do science, namely conjectures, ideally in the form of a mathematical model, and refutations"; the two added dimensions are "the amount of attention paid to the observer and the amount of impact of a theory on the system described"; and "to accept methods for acting as well as theories as a contribution to science." The full text was not read.

[9] Ranulph Glanville, "Try Again. Fail Again. Fail Better," *Kybernetes* 36, no. 9/10 (2007): 1173–1206, for the design conversation (**unverified**, full text not obtained); Ranulph Glanville, "Researching Design and Designing Research," *Design Issues* 15, no. 2 (1999): 80–91, for the second point: "it is inappropriate to require design to be 'scientific': for scientific research is a subset (a restricted form) of design" (read 29 September 2026 in the author's typescript, "Re-searching Design and Designing Research," https://constructivist.info/radical/papers/glanville/glanville98-design.pdf; the *Design Issues* page has not been checked). Glanville elsewhere calls the VSM the system that "epitomizes the truly second-order cybernetic system" ("Second Order Cybernetics," in *Systems Science and Cybernetics*, ed. F. Parra-Luna, vol. 3, Encyclopedia of Life Support Systems, EOLSS Publishers; the year of the online edition, 2002 or 2003, is still to be confirmed). No second-order critique of the VSM has been located, and that absence has not been searched systematically.

[10] Ernst von Glasersfeld, "Why Constructivism Must Be Radical," in *Constructivism and Education*, ed. Marie Larochelle, Nadine Bednarz and Jim Garrison (Cambridge University Press, 1998), 23–28: "an action, operation, conceptual structure, or even a theory, is considered 'viable' as long as it is useful in accomplishing a task or in achieving a goal that one has set for oneself." Wording from the von Glasersfeld archive copy (https://www.evg-archive.net/), which describes itself as an approximate translation of a French original; the published CUP wording has not been checked. The equivocation argument that follows is this book's own. A search on 23 September 2026 (von Glasersfeld with viability and Beer or the VSM) found no source that treats it. It should be put to constructivist readers before print, and presented as an argument rather than as a finding.

[11] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 1, quoting the *Oxford English Dictionary* for "able to maintain a separate existence."

not necessarily agree on the existence, or nature, or boundaries of any system so detected"; and "it is possible to say only that the treatment of a certain collection of things as a system is helpful," a point to be tested "by checking on the predictive capability of the insight."[12] Which states of the system count depends on "the purposes of the system."[13] For Beer, then, viability is a property of a system that an observer has recognized for a purpose. For von Glasersfeld it is a property of the knowledge. Two concepts wear one word, and it is the word this tradition is named for. They are nearer to each other than a quick reading suggests, and part of the distance between them runs through Beer's own pages.

The consequence runs in one direction only, which makes it dangerous and not merely confusing. Slide from Beer's sense to von Glasersfeld's, and "the model is viable" stops meaning that the model is right about what keeps organizations alive and starts meaning that the model has been useful to the people using it. That second sentence is unfalsifiable. It is also, near enough, the sentence the field actually says. Beer's account makes the slide easy, because it grounds the recognition of a system in its being helpful. It also supplies the brake, a test of predictive capability, and the brake is the part the field does not use. Here the slide gets its name; Chapter 23 will fix the narrow claim in Beer's sense precisely so that it cannot happen unnoticed.

We have found no source that treats this equivocation, so the argument is ours and must be presented as ours. It may also be wrong, and the way to find out is to put it in front of the people whose word it is.

What this chapter does not establish

It does not establish that the ASC has failed at anything. The society has never claimed to be an evidence-producing body, and holding it to a standard it does not profess would be the error this book accuses others of making.

Nor does it establish that second-order cybernetics is an obstacle. Its central claim, that the observer belongs in the account, is one this book accepts and one Chapter 29 is built to take seriously.

What it establishes is that the discipline's official self-description offers complete shelter from testing, that its incentive structure does not reward refutation, and that the tradition's central term has two senses in active circulation, and the field slides between them. None of these is anybody's fault. All three are ours to work around, because we are the ones making claims about organizations rather than about ways of thinking.

[12] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 242 for "defined by intelligence," 243 for the other quotations.

[13] Beer, *Diagnosing the System*, 98: "What matters in discriminating one 'state of the system' from other states is whether the resulting change of state serves, or has no bearing on, the purposes of the system." On the same page Beer glosses the purpose implied by viability, to survive, as "the 'separate existence' of the 'recognizable entity'."

Receipt — say which layer each claim is on, in the teaching material. The target is specific and already public: SCiO's five-course VSM sequence (VSM1 to VSM5, taught by Patrick Hoverstadt; the first four sold pre-recorded at £200 each, the fifth on enquiry), Metaphorum's public "Viable System Model" page (Chapter 17), the diagrams reproduced in them, and the two paragraphs on SCiO's own "Why use the Viable System Model?" page, where "has proven to be a scalable, super-fast, flexible, rigorous, approach" and "a systemic model of what any type of organisation needs in order to be able to survive in a changing environment" sit within a few lines of each other — the first an instrumental claim, the second a constitutive one, and nothing in the page marking the difference.[14] Label every claim in that set as one of three kinds: a way of thinking (no failure condition, and none claimed), an instrumental claim (assessed by usefulness, with the measure named), or a constitutive claim (assessed by observation, with the disconfirming observation named, per Chapter 28). Publish the labelled set. Price: two working days per course module, done once, by the people who already own the material; a full pass over one certification syllabus is inside a week. The result is falsifiable in a modest but real sense: if the constitutive column comes out empty, then the model is a way of thinking and the neuroanatomical mapping, the necessary-and-sufficient claim and the pathology derivations must leave the teaching, which is a substantial and specific loss. If the column is not empty, every entry in it is now a debt with a named test, and Chapter 30 can price them. What the exercise cannot produce is the present situation, in which the three kinds of claim are taught in one breath and defended by whichever standard the question requires.

Handover

The shelter described here is institutional. The next chapter takes the version that arrives as method, from the neighbour best equipped to give it.

[14] SCiO, "Why use the Viable System Model (VSM)?" and the courses catalogue, https://www.systemspractice.org/why-practice-systems-thinking/vsm and https://www.systemspractice.org/courses. Checked 14 and 29 September 2026 for the quoted phrases, the five VSM courses and the prices (VSM1–VSM4 "£200 pre-recorded," VSM5 "please enquire"). The catalogue states that its tutors have "collectively trained over ten thousand people in systems thinking to masters level" and "written 23 books on systems thinking and practice"; if the first figure holds, it is the scale at which the unlabelled claims have travelled.

Chapter 20

Usefulness Instead of Truth

The claim

On the second page of *Diagnosing the System for Organizations* (1985), Beer sets three short lines apart from his text: "So remember: / a model is neither true nor false: / it is more or less useful."[1] It is the most respectable sentence in the systems literature, and it is not wrong. It is also the most damaging sentence in our inheritance, and we did not borrow it. It is the founder's, and it was not new in 1985. In *Decision and Control* (1966) a methodology is "neither right nor wrong, but only more or less useful," and later in *Diagnosing the System* the model has no "correct" interpretation, only "more or less useful interpretations."[2]

In 1966 Beer also said how the usefulness of a model is to be measured: "by straightforward testing of the predictive value of the scientific model. If this predictive value is low, we had better start again."[3] The sentence came with a way of losing. The tradition kept the sentence and dropped the test.

Our neighbors reached the same sentence by their own routes, and the most careful of them built the machinery that makes it safe. The machinery is the subject here, and so is what the sentence becomes when it travels without it.

[1] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2. The lines follow "you are not determining absolute facts: you are establishing a set of conventions," in the passage on fixing the organization to be modelled and its boundaries.

[2] Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 96, where "a cohesive collection of fundamental methods, which may be called a methodology, provides a framework for research which is neither right nor wrong, but only more or less useful"; Beer, *Diagnosing the System*, 144: "There is no 'correct' interpretation of the VSM. We have spoken instead of more or less useful interpretations. Even so, there may well be incorrect interpretations, in the sense that the model's power to account for viability may become denatured by their use."

[3] Beer, *Decision and Control*, 113. The passage describes the scientific model in operational research in general, not the Viable System Model in particular; Chapter 21 quotes it at more length. Jackson reads Beer's 1966 account the same way: the cybernetic concepts were to turn management's conceptual models into a homomorphic model "formulated rigorously enough to be falsifiable, and so scientific in Popper's terms," and a valid model "can be tested for its predictive value"; he adds that in Beer's view "the model is yet to be falsified" (Michael C. Jackson, *Critical Systems Thinking and the Management of Complexity* (Wiley, 2019), 299–300, checked 29 September 2026 in the calibre PDF).

The neighbour who earned it

System Dynamics has the most developed validation philosophy of any tradition adjacent to ours, and it arrived at the usefulness position by argument, not by convenience.

Forrester and Senge state that there is no method for proving a model correct, and that confidence rather than proof is the proper criterion.[4] Barlas and Carpenter engage Popper directly, including the Duhem problem (a failed prediction never tells you which of your assumptions to give up), and conclude that models lie on a continuum of usefulness.[5] Barlas gives the operational form: right behavior for the right reasons.[6] Sterman rejects validation in the sense of establishing truthfulness.[7]

So far this is our sentence, said better. The next part is the one we did not take.

Having given up proof, System Dynamics built structured test batteries in its place. Barlas separates structure tests from behavior tests. The first ask whether the model's mechanisms correspond to the mechanisms in the system, independently of whether the output looks right; the second ask whether the output matches, judged against patterns and not point predictions. Forrester and Senge had already named them in 1979: structure-verification, parameter-verification, extreme-conditions, boundary-adequacy and dimensional-consistency tests of structure, and behavior-reproduction, behavior-anomaly and behavior-sensitivity tests of behavior, among others, each with a stated procedure.[8] A model that fails them is in trouble, however useful anyone found it.

"More or less useful" in that tradition is the conclusion of a program with named ways of losing. In ours it began with one, Beer's prediction test. Beer stated the test; neither he nor anyone after him turned it into a battery

---

[4] Jay W. Forrester and Peter M. Senge, "Tests for Building Confidence in System Dynamics Models," in *System Dynamics*, ed. Augusto A. Legasto Jr., Jay W. Forrester and James M. Lyneis, TIMS Studies in the Management Sciences 14 (North-Holland, 1980), 209–28. Read in the MIT System Dynamics Group memo D-2926-7 of 8 June 1979, 8: "We believe confidence is the proper criterion because there can be no proof of the absolute correctness with which a model represents reality. There is no method for proving a model to be correct." The page in the published version has not been checked.

[5] Yaman Barlas and Stanley Carpenter, "Philosophical Roots of Model Validation: Two Paradigms," *System Dynamics Review* 6, no. 2 (1990): 148–66. **Unverified.**

[6] Yaman Barlas, "Formal Aspects of Model Validity and Validation in System Dynamics," *System Dynamics Review* 12, no. 3 (1996): 183–210. **Unverified:** the article has not been read; its separation of structure tests from behavior tests is carried from the research note. The list of named tests in the text is taken from Forrester and Senge (see the note to that sentence), not from Barlas.

[7] John D. Sterman, "All Models Are Wrong: Reflections on Becoming a Systems Scientist," *System Dynamics Review* 18, no. 4 (2002): 501–31. **Unverified.**

[8] Forrester and Senge, memo D-2926-7, 2 (table of contents): section III, "Tests of Model Structure" (structure-verification, parameter-verification, extreme-conditions, boundary-adequacy and dimensional-consistency tests); section IV, "Tests of Model Behavior" (behavior-reproduction, behavior-prediction, behavior-anomaly, family-member, surprise-behavior, extreme-policy, boundary-adequacy and behavior-sensitivity tests); section V, tests of policy implications. Each test is described in its own subsection.

for the Viable System Model, with named procedures and a threshold for "low." We kept the epistemology, skipped the battery, let the founder's test lapse, and were left with a sentence that sounds like methodological sophistication and works as an exemption.

The right response is to pay for the position, and the structure/behavior split is directly portable. A VSM diagnosis makes structural claims (this function is present, this channel carries this variety, this unit is a System One at this level) and behavioral claims (this pathology will produce this oscillation). Barlas's distinction says the two must be tested separately, and that a diagnosis passing the second while failing the first is not corroborated. The node, edge and subgraph typing in the companion working paper, *A Stress Test for the Viable System Model*, connects to this directly,[9] and Schwaninger works in both traditions, so the bridge is a short one to build.

The same move, three more times

Replacing the falsification criterion is the field's standard maneuver, well beyond System Dynamics, and almost nobody in the neighborhood rejects testing outright.

The Hull school judges a methodology by fit to the problem context, by pluralism, and by practical and emancipatory results, not as a hypothesis. Michael Jackson's position is that systems thinking should learn which approaches bring about beneficial change in a particular context; Gerald Midgley holds that theories are more or less useful depending on the purposes of intervention.[10] Within that frame the VSM is the functionalist lens, and the metamethodology does not ask whether it is *right* about organizations. Jackson himself does, from the functionalist side: an empirical research agenda around the VSM is, he writes, "certainly essential, and it is disappointing that so little has been done."[11]

---

[9] Krishan Mathis and Margeret Heath, "A Stress Test for the Viable System Model — A Cumulative Research Program," open preprint (Zenodo, 2026), https://doi.org/10.5281/zenodo.22306416, §2.5, where claims are typed as node, edge and subgraph claims.

[10] Michael C. Jackson, "Pragmatism and Critical Systems Thinking: Back to the Future of Systems Thinking," *Integration and Implementation Insights* (i2Insights), 4 April 2023, https://i2insights.org/2023/04/04/pragmatism-and-critical-systems-thinking/; Gerald Midgley, "Theoretical Pluralism in Systemic Action Research," *Systemic Practice and Action Research* 24, no. 1 (2011): 1–15, https://doi.org/10.1007/s11213-010-9176-2, where theories "are more or less useful depending on the purposes of intervention" (checked against the abstract). Jackson's critique of the VSM is in *Critical Systems Thinking and the Management of Complexity* (Wiley, 2019), §13.4, pp. 325–37 (checked 23 September 2026 against the calibre copy of the 2019 Wiley edition). He classes the VSM as "functionalist in orientation," though structuralist rather than positivist (p. 329), and places it in the complex-unitary cell of his system of systems methodologies (pp. 334–36). Chapter 16 places Jackson's argument.

[11] Jackson, *Critical Systems Thinking and the Management of Complexity*, 330, checked 29 September 2026 in the calibre PDF. The agenda he has in mind concerns "what structures are beneficial for viability and which impair it"; he raises it while discussing Achterbergh and Vriens (2011).

Werner Ulrich replaces the criterion with boundary critique, deliberately and with an argument: his charge against critical rationalism is that it has not really overcome the empirical foundationalism of earlier positivist research practice.[12] Chapter 16 takes his objection at full strength and concedes what it establishes. The concern here is narrower. Boundary critique is a criterion for the legitimacy of an intervention, and it cannot settle whether there are five subsystems. An empirical program does not answer Ulrich, Ulrich does not answer the empirical question, and this book should say both.

Peter Checkland is the one who ought to be named an ally rather than a false friend. He and Holwell accept that testing hypotheses to destruction carries the strongest truth claim in natural science, and for social inquiry they substitute recoverability: a research process based on a prior declaration of the epistemology, such that another person can retrace what was done and why.[13] Prior declaration is exactly what the receipt protocol asks for. Recoverability is a weaker obligation than a test, but a real one. This book can adopt it without amendment, and it is the model for what the other replacements should look like.

Across all four the pattern is the same. Each substitute criterion is adequate for instrumental claims: whether a diagnosis helps, whether an intervention is legitimate, whether an account can be retraced. None can decide a constitutive claim: that the five functions are necessary, that they are sufficient, that System Five corresponds to a cortex. Those claims are in our teaching and in our books. Usefulness cannot assess them, because a false model can be useful, and frequently is.

Beer's 1966 account shows how a constitutive claim escapes. A scientific model, he wrote, "actually defines the extent to which [two situations] are structurally identical," and if it has ignored as irrelevant something that is in fact relevant, "the model will lose in utility, but it cannot lose in validity. This is its peculiar strength."[14] The claim of structural identity is secured by construction, and prediction is left as its only exposure. Drop the prediction test, and nothing is exposed.

---

[12] Werner Ulrich, "Rethinking Critically Reflective Research Practice: Beyond Popper's Critical Rationalism," *Journal of Research Practice* 2, no. 2 (2006): P1. The charge that critical rationalism has "not really overcome the empirical foundationalism" of earlier positivist research practice is checked against the abstract; the full text is open (ERIC EJ805693) but has not been read.

[13] Peter Checkland and Sue Holwell, "Action Research: Its Nature and Validity," *Systemic Practice and Action Research* 11, no. 1 (1998): 9–21. **Unverified.**

[14] Beer, *Decision and Control*, 113. The passage describes the scientific model in operational research in general, not the Viable System Model in particular; Chapter 21 quotes it at more length. Jackson reads Beer's 1966 account the same way: the cybernetic concepts were to turn management's conceptual models into a homomorphic model "formulated rigorously enough to be falsifiable, and so scientific in Popper's terms," and a valid model "can be tested for its predictive value"; he adds that in Beer's view "the model is yet to be falsified" (Michael C. Jackson, *Critical Systems Thinking and the Management of Complexity* (Wiley, 2019), 299–300, checked 29 September 2026 in the calibre PDF).

The asymmetry that makes this a friendship

A complaint that the neighbors refuse to test anything would be easy to write. It would also be false, and it would miss what is actually happening.

System Dynamics runs test batteries. The Hull school evaluates interventions with stakeholders. Researchers on problem-structuring methods have built evaluation frameworks and are candid that claims for these methods usually rest on single case studies of intervention, and that such modeling yields insights rather than testable results.[15] That candor is more than this tradition has managed. Ormerod maintains a critical-rationalist strand inside operational research and is probably the closest ally the program has in that field.[16]

Every one of these communities tests its own work more than we test ours, while holding an epistemology at least as permissive. The philosophy is shared. They treated it as an obligation to build something else; we treated it as permission to build nothing.

That is why they are false friends and not adversaries, and why the word "false" does narrow work. Nothing they hold threatens this model. What they hold has been used, by us, as a reason not to find out whether the model is any good. Their versions confirmed a sentence we already had from Beer, and a position confirmed by respected neighbors does not feel like a choice that needs defending.

What this chapter does not establish

It does not establish that usefulness is the wrong criterion for instrumental claims. It is the right one, and Chapter 24 will assess the model's instrumental claims by exactly that standard.

Nor does it establish that System Dynamics, the Hull school or critical systems heuristics owe this tradition anything. They owe us nothing. The sentence came from our own founder, and so did the test we dropped.

---

[15] Gerald Midgley et al., "Towards a New Framework for Evaluating Systemic Problem Structuring Methods," *European Journal of Operational Research* 229, no. 1 (2013): 143–54, https://doi.org/10.1016/j.ejor.2013.01.047, for the reliance on "single case studies of intervention"; Fran Ackermann, "Problem Structuring Methods 'in the Dock': Arguing the Case for Soft OR," *European Journal of Operational Research* 219, no. 3 (2012): 652–58, https://doi.org/10.1016/j.ejor.2011.11.014, for "insights rather than testable results." Both checked against the abstracts only.

[16] Richard J. Ormerod, "The History and Ideas of Critical Rationalism: The Philosophy of Karl Popper and Its Implications for OR," *Journal of the Operational Research Society* 60, no. 4 (2009): 441–60, https://doi.org/10.1057/palgrave.jors.2602573; Richard J. Ormerod, "Critical Rationalism for Practice and Its Relationship to Critical Systems Thinking," *Systems Research and Behavioral Science* 33, no. 1 (2016): 4–23 (online 2014), https://doi.org/10.1002/sres.2326. Identified from the metadata and abstracts; the full texts are paywalled and have not been read.

What it establishes is that the tradition took a criterion suited to one class of claim and applied it to another, that it dropped the test its founder attached to the criterion, and that System Dynamics shows what the criterion costs when taken seriously.

Receipt — run one diagnosis through the structure/behaviour battery. Take a single completed and published VSM diagnosis, for example the M□Clean case in Espinosa's 2023 book,[17] and subject it to the System Dynamics test set in Barlas's form, adapted minimally: boundary adequacy (does the recursion cut include what drives the behaviour at issue?), structure verification (is each claimed function and channel independently confirmable by someone who did not produce the diagnosis?), parameter and structure sensitivity (does the diagnosis change if a defensible alternative recursion cut is taken?), extreme-condition behaviour (does the model's predicted pathology appear when the named variety imbalance is driven to its limit?), and behaviour-pattern reproduction (did the predicted oscillation occur?). On a published case the extreme-condition item runs on the case record: find the period of the largest named variety imbalance and check whether the predicted pathology appeared then. Publish the battery and the result together. Price: one analyst for two to three weeks on an existing written-up case, plus a second analyst for three days on the independent structure verification; the sensitivity item is desk work. If a published diagnosis cannot survive structure verification by an independent analyst, the field learns that its diagnoses are not reproducible, which is Chapter 5's charge made concrete and would be the most consequential negative result available at this price. If it survives, the field has its first diagnosis with a battery attached, and a template the next one can be held to.

Handover

Three chapters have taken shelters made of criteria. The last of them is made of a distinction, and it is the strongest.

[17] Angela Espinosa, *Sustainable Self-Governance in Businesses and Society: The Viable System Model in Action* (Routledge, 2023), 116–17, "Case Study – M□Clean (with A.C. Martinez)": a six-month project in a Mexican cleaning-services corporation with six regional offices, with a recursive analysis (fig. 4.4) and a VSM diagnosis of a small service unit (fig. 4.5). Checked 29 September 2026 in the calibre copy of the first edition. Whether the published account records enough of the case to run every item of the battery has not been checked.

Chapter 21

The Category Error as Shelter

The claim

Ask for a test of the Viable System Model, and the strongest reply this tradition has is not that the model has passed one. It is that the request mistakes what kind of thing the model is.

That defense is good. It has a serious literature behind it, people with no stake in the VSM hold it, and this book has not answered it. It deserves to be stated at full strength, priced, and set beside a neighboring tradition that made the same argument in public and paid what it costs.

Pickering's version

Andrew Pickering's history of cybernetics makes the argument for the whole field. On his account, cybernetics is about a shift from epistemology to ontology, from representation to performativity, and Beer's work in particular displays a cybernetic ontology of unknowability and becoming.[1]

Read that carefully, because it is no soft position. It does not say that the model is hard to test, or that testing it is premature. It says that the model was never a representation of an organization in the first place: it is a way of engaging with a system whose behavior cannot be known in advance, a set of moves rather than a set of assertions. On that reading, asking whether there really are five subsystems is like asking whether a hammer is true.

Pickering has the historical evidence for it. Beer's practice was performative in exactly this sense: the biological computers, the Chilean operations room, the insistence that the model is built in the act of using it, the refusal to treat organizations as objects sitting still to be described. A reader who knows the archive will recognize Beer's practice in Pickering's account. Beer's books say something more definite about what the model is, in his

[1]Andrew Pickering, "Cybernetics and the Mangle: Ashby, Beer and Pask," *Social Studies of Science* 32, no. 3 (2002): 413–37, https://doi.org/10.1177/0306312702032003003; Andrew Pickering, "The Science of the Unknowable: Stafford Beer's Cybernetic Informatics," *Kybernetes* 33, no. 3/4 (2004): 499–521, https://doi.org/10.1108/03684920410523535, whose abstract speaks of a "cybernetic ontology of unknowability and becoming." **Unverified:** the phrase "from epistemology to ontology, from representation to performativity" is carried from a research note and is not in the 2002 abstract; neither full text has been read. Obtain the papers, quote exactly, and — since this chapter states his position at full strength — check that the strengthening is his and not ours.

own words. Even in his most practice-minded pages he asks that models be open to checking: the model of the whole enterprise that senior managers carry in their heads fails, he writes, because "it is not made explicit, and cannot be interrogated or tested."[2]

If Pickering is right, better measurement does not help. It refines an instrument aimed at nothing, the structure Chapter 16 identified in the objections that land.

Beer's own declaration

Beer did not leave the question to his readers. In *Decision and Control* (1966) he asked what kind of thing a model is and answered with two refusals. Can it be equated with a scientific hypothesis? "Clearly not; the model does not postulate the causal mechanisms that underlie events, it simply represents the pattern of the events themselves in advance". Is it a scientific theory, then? "Again, it is not; it has no explanatory content. A model is simply a reflection of whatever is the case, which is explicitly made available for experimentation."[3]

Twelve pages later he said what that status buys. A scientific model "is a homomorphism on to which two different situations are mapped, and which actually defines the extent to which they are structurally identical." If the mapping leaves out something that matters, "the model will lose in utility, but it cannot lose in validity. This is its peculiar strength." Utility, in turn, is measured "by straightforward testing of the predictive value of the scientific model. If this predictive value is low, we had better start again."[4] Nineteen years later the same status reached practitioners as a maxim: a model "is neither true nor false: it is more or less useful."[5]

This is a declaration in two layers, made by the founder, in print. It is neither Pickering's reading nor a hypothesis. Validity is secured by construction; utility is exposed to prediction. And the representational claim stays in. A model that "defines the extent to which [two situations] are structurally identical" says something about the organization, and the declaration protects that something from being wrong. The performative reading would give the claim up. Beer's declaration keeps it and sets a single test beside it. Beer held the other reading as well: closing the chapter of the same book on controlling enterprises, whose model is the precursor of the Viable System Model, he called that model "no more and no less than a theory" and "a descriptive account of how enterprises actually work", so

[2]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 237, on System Four's "contained model of the whole viable system" held "in the heads of the senior people." Checked 29 September 2026 in the calibre PDF.

[3]Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 101.

[4]Beer, *Decision and Control*, 113. The passage describes the scientific model in operational research in general; Chapter 20 uses the same page.

[5]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2, set as three lines: "So remember: / a model is neither true nor false: / it is more or less useful."

the founder stood on both readings, one that can be wrong and one that "cannot lose in validity".[6]

What it would cost

The book's answer is to price the performative reading, and the price is high enough that nobody in this tradition has been willing to pay it in public.

If the VSM is performative rather than representational, its representational claims must leave the teaching. All of them.

The neuroanatomical mapping goes. If the model is a way of acting, System Five corresponds to nothing in a cortex, and Chapter 1's charge about what Beer actually wrote becomes moot, because the mapping was never making a claim. The necessary-and-sufficient claim goes: a set of moves is not the kind of thing that could be necessary and sufficient for the survival of an organization. Beer's own presentation of *The Heart of Enterprise* as written in the belief that the necessary and sufficient conditions of viability had been established goes with it, though he had warned at the outset that the rigour would "fall short of the target."[7] The derived pathology catalogue survives only as a repertoire of things to try, no longer as predictions. Beer's comparison of Ashby's Law to the law of gravity goes,[8] and so does every diagram taught as a picture of how organizations are.

The subtraction is large: most of what is currently taught, and all of what is currently claimed for the model when it is sold.

A community is entitled to take the performative reading. It is not entitled to take the performative reading when asked for evidence and the representational reading when teaching, selling or citing. That alternation is no philosophical position, and it is what makes the model unchargeable.

---

[6] Beer, *Decision and Control*, 397, at the end of chapter 15, "Controlling Enterprises", under "a special declaration": "This is the only case, of the very large number of studies and models and approaches put forward in this book, in which what is no more and no less than a theory has been advanced." The phrase "a descriptive account of how enterprises actually work" follows in the same paragraph. Pages follow the Stafford Beer Classic Library reprint (John Wiley & Sons, 1994), as for pages 101 and 113.

[7] Stafford Beer, *The Heart of Enterprise* (Chichester: John Wiley & Sons, 1979), 408: "But if a system has all of the necessary and sufficient conditions of viability, which we set out to discover and have now listed, it will live." The method is announced at p. 115 — "our method will be to aim for a statement of conditions that are necessary and sufficient" — with a reservation: "As a good philosopher, I am bound to declare the inadequacy of the rigour; and that means that I shall fall short of the target." Checked 23 September 2026 against the calibre copy of the 1979 Wiley edition and re-checked 29 September 2026. The Part Three opener quotes the same passage from p. 408; Chapter 9 keeps p. 235 for the phrase "the necessary and sufficient conditions of any viable system."

[8] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 89: "I consider that this law stands in the same relation to management as the law of gravity stands to Newtonian physics"; Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 201: "I think that Ashby's Law of Requisite Variety is the managerial Law of Gravity." Both checked 29 September 2026 in the calibre PDFs.

The price named above is also a revision rule, of the form Chapter 28 asks for: if the performative reading is adopted, these specific claims are withdrawn from the teaching material by a stated date. It costs nothing to state, and nobody has ever stated it.

The tradition that did it in public

One tradition in the cybernetic lineage has made exactly this argument, in print, about its own central construct, and accepted what followed.

Asked directly whether the free energy principle can be falsified, Karl Friston answers that the principle is what it is: a principle; that like Hamilton's Principle of Stationary Action it cannot be falsified and cannot be disproven; and that hypotheses about the brain performing some form of Bayesian inference or predictive coding are what they are: hypotheses.[9] Chapter 9 quotes the same passage for a different purpose. What matters here is its shape: an explicit two-layer declaration, made by the author, unprompted, in print.

The philosophical response divides three ways, and the division is instructive. Mel Andrews argues that demanding falsifiability of the principle is a category error, which is Pickering's move made analytically.[10] Matteo Colombo and Cory Wright hold that the epistemic status of the principle is unclear.[11] Jelle Bruineberg and colleagues argue that its metaphysical claims cannot be justified by appeal to the success of the mathematical framework alone.[12] Outside that debate, Anil Seth, placing predictive processing in the cybernetic line, calls the supporting evidence only circumstantial.[13]

Look at what the declaration bought and what it cost. It bought the principle immunity from a demand that did not apply to it. It cost the community the right to claim empirical support for the principle from the success of the process theories, and it exposed the process theories to

[9] Karl Friston, Martin Fortier and Daniel A. Friedman, "Of Woodlice and Men: A Bayesian Account of Cognition, Life and Consciousness; An Interview with Karl Friston," *ALIUS Bulletin* 2 (2018): 17–43, at 21. The passage, as quoted in Chapter 9 from the same page: "the free energy principle is what it is—a principle. Like Hamilton's Principle of Stationary Action, it cannot be falsified. It cannot be disproven." The paraphrase above follows it.

[10] Mel Andrews, "The Math Is Not the Territory: Navigating the Free Energy Principle," *Biology & Philosophy* 36 (2021): 30. **Unverified:** characterization from an abstract-level note.

[11] Matteo Colombo and Cory Wright, "First Principles in the Life Sciences: The Free-Energy Principle, Organicism, and Mechanism," *Synthese* 198 (2021): 3463–88. **Unverified.**

[12] Jelle Bruineberg et al., "The Emperor's New Markov Blankets," *Behavioral and Brain Sciences* 45 (2022): e183, https://doi.org/10.1017/S0140525X21002351. The paraphrase follows the published abstract: Friston blankets require "additional philosophical premises and cannot be justified by an appeal to the success of the mathematical framework alone"; the 2020 PhilSci-Archive preprint (18467) argues the same at length.

[13] Anil K. Seth, "The Cybernetic Bayesian Brain: From Interoceptive Inference to Sensorimotor Contingencies," in *Open MIND*, ed. Thomas Metzinger and Jennifer M. Windt (MIND Group, 2015), 35(T), https://doi.org/10.15502/9783958570108, 4: the literature "so far provides only circumstantial support for the basics of PP."

ordinary testing, where their evidence base is still described as limited.[14] Immunity and exposure came as a pair. That is what an honest category-error defense looks like.

Set the two founders side by side. Friston put the principle beyond testing and left the process theories exposed. Beer put the model's validity beyond testing and left its utility exposed to prediction. Both declarations pair an immunity with an exposure, and both were made in print. The difference lies in what the communities did next. Friston's took the exposure with the immunity, and its process theories are being tested. Ours kept the immunity, left the prediction test unbuilt, and invokes the immunity only when asked.

The shelter has a version one floor down, in the formal layer, and Chapter 3 has already dealt with its substance. The model borrows authority from the Conant–Ashby theorem,[15] and since a theorem makes no empirical hypothesis, that layer too is immune by construction. Beer came close to saying so for Ashby's law, in the reply on tautology that Chapter 9 quotes. What has changed since Chapter 3 was drafted is that the theorem's own standing is now under formal attack, and the attack comes from outside our field. Chapter 22 takes it up, because it is the one place where the adversaries are doing us a favor.

What this chapter does not establish

It does not establish that Pickering is wrong. This book has no argument that he is, and the honest statement of our position is that the performative reading has been neither refuted nor accepted.

Nor does it establish that the free energy principle community is a model in every respect. Its declaration has been criticized from three directions, and the criticisms are live.

What it establishes is that this tradition holds a defense it has never declared, the performative reading, beside one its founder did declare and it has not honored, the prediction test attached to Beer's own status for the model. Declaring the first carries a price nobody has quoted; honoring the second carries one nobody has paid. A community that wants either defense without its price is doing something other than philosophy.

Receipt — the two-layer declaration, with the subtraction listed. Publish a two-page statement, signed by the certifying bodies, dividing the Viable System Model into a principle layer and a hypothesis layer on the

[14] Kevin S. Walsh, David P. McGovern, Andy Clark and Redmond G. O'Connell, "Evaluating the Neurophysiological Evidence for Predictive Processing as a Model of Perception," *Annals of the New York Academy of Sciences* 1464 (2020): 242–68, https://doi.org/10.1111/nyas.14321, whose abstract says predictive processing models "are also criticized for lacking the empirical support to justify their status" and describes a "limited evidence base" and a "recent surge" of research to fill it (checked against the abstract); Seth, "Cybernetic Bayesian Brain," 4.

[15] Roger C. Conant and W. Ross Ashby, "Every Good Regulator of a System Must Be a Model of That System," *International Journal of Systems Science* 1, no. 2 (1970): 89–97, https://doi.org/10.1080/00207727008920220. Chapter 3 discusses what the theorem proves.

template Friston set for the free energy principle, or on the one Beer set in 1966, in which validity is secured by construction and utility is tested by prediction: what the model asserts that no observation could contradict, and what it asserts that observation could. Attach to the principle layer the list of claims that must consequently be withdrawn from teaching and from client-facing material — the neuroanatomical mapping, the necessary-and-sufficient formulation, the Ashby–Newton comparison, and any prediction currently made from the pathology catalogue — with a date by which the withdrawal takes effect. Price: one workshop day to draft, one round of comment, one signature; the expensive part is not the writing but the consent, and the consent is the deliverable. The exercise is decisive in a way the book's other receipts are not: it cannot be completed without someone giving something up, and if it cannot be completed at all, that is the strongest available evidence for Chapter 9's central charge — not that the tradition has failed a test, but that it will not enter a position from which anything could be lost.

Handover

Five chapters have dealt with friends. The last deals with the people who are supposed to be the danger, and asks whether any of them is.

Chapter 22

The Enemies We Actually Have

The claim

Count the people actually attacking the Viable System Model, and the list is short. This tradition believes itself embattled. It is unattended, which is worse.

The standard is Chapter 16's, unchanged. An objection lands when it identifies something the model needs in order to do the job it claims, cannot supply from its own resources, and has not been given by any amendment since. By that standard three objections land, Ulrich's, Hayek's and Cooper's, and Chapter 16 has already taken them. Nothing here adds to that list. What remains is everything else commonly treated as opposition, sorted by what it costs us.

The rival ontology

Dave Snowden is the most visible living critic of this model in the market where most of its practitioners work. His position, stated in December 2025, distinguishes Bateson-style cybernetics from its use for the VSM and its relatives, and holds that had Beer known complexity science he would have created something other than the VSM.[1]

Taken seriously, the claim is an ontological objection: much of organizational life is complex in Cynefin's sense, where "cause and effect are only coherent in retrospect and do not repeat," and a model that maps functions and channels is aimed at a kind of object that is not there.

What does it cost us? Commercially, a great deal. Cynefin is embedded in the same DACH consulting market this tradition sells into, and a

[1] Dave Snowden, "Positioning Systems Approaches," The Cynefin Co blog, 24 December 2025, https://thecynefin.co/positioning-systems-approaches/, accessed 29 September 2026: "I also read and fell in love with much of Beer's work, but still maintain that had he known about complexity science, he would have created something other than VSM"; "I think we need to make a rigid distinction between how cybernetics, a la Bateson, is used and how it is used for VSM et al." Snowden presents the post as a proposal he is "not fully committed to" yet. The ontological reading in the following paragraph follows the framework's founding statement: C. F. Kurtz and D. J. Snowden, "The New Dynamics of Strategy: Sense-Making in a Complex and Complicated World," *IBM Systems Journal* 42, no. 3 (2003): 462–83, where in the complex domain "cause and effect are only coherent in retrospect and do not repeat" (fig. 1, p. 468; "retrospective coherence," p. 469) and systems thinking is placed in the ordered, knowable domain, "too often confused with complexity theory" (p. 468). Checked 23 September 2026 against the calibre copy of the journal article.

practitioner who loses that argument in a room loses the engagement. Epistemically, nothing yet. No test is offered. Cynefin's authors say so themselves: its "value is not so much in logical arguments or empirical verifications as in its effect on the sense-making and decision-making capabilities of those who use it."[2] There is no observation that would settle whether organizations are the kind of object the VSM diagnoses or the kind Cynefin describes, and neither side has proposed one. Two ontologies facing each other without a discriminating experiment: that is the condition this book charges the VSM with, now held by both parties.

By the standard, it does not land. It is still the most interesting place to look for an adversary who could be made to land, and the way to make that happen is to propose the experiment ourselves.

The sweeping dismissals

An older literature rejects systems thinking wholesale. Ida Hoos holds that systems analysis is impervious to criticism; David Berlinski's judgment is that the subject is a sham; Robert Lilienfeld treats systems theory as ideology.[3]

Our field usually cites these as evidence of hostility. They are evidence of something else. None of them is about the Viable System Model. They address systems analysis and general systems theory in the 1970s, they predate most of what this tradition has produced, and their charge, that the subject is untestable and immune to criticism, is one this book makes in more detail and with more evidence than any of them assembled.

They cost us a reputation more than an argument, and in the specific respect they name, the reputation is deserved. The right response is to note that the general charge was made fifty years ago, that this tradition never answered it, and that the answer was always going to have to take the form of a test, never of a reply.

They do not land, for the pedestrian reason that they do not engage the model. Nor should they serve as a foil: a tradition that defends itself by pointing at its weakest critics is doing what Chapter 17 describes.

Friendly fire on the foundation

The most substantial recent criticism of anything this tradition rests on comes from people who have never heard of Beer.

---

[2] Kurtz and Snowden, "The New Dynamics of Strategy," 468: "We consider Cynefin a sense-making framework, which means that its value is not so much in logical arguments or empirical verifications as in its effect on the sense-making and decision-making capabilities of those who use it." Checked 29 September 2026 in the calibre copy, as was the figure quoted in the paragraph above (fig. 1, 468; see the note on Snowden).

[3] Ida R. Hoos, *Systems Analysis in Public Policy* (1972), on systems analysis as impervious to criticism — **recorded from a secondary source and unverified.** David Berlinski, *On Systems Analysis* (1976/1978): "the subject is a sham" — **unverified.** Robert Lilienfeld, *The Rise of Systems Theory* (1978) — **unverified, secondary.** Any Popper statement on cybernetics: not found.

The Conant–Ashby theorem, that every good regulator of a system must be a model of that system, supplies the formal authority VSM teaching borrows when it wants to be more than a heuristic. Chapter 3 deals with what the theorem says and what we quote instead. Since then the theorem's own standing has come under attack. Iqbal Thobani argues in *Synthese* that the notion of an internal model in these principles is far too permissive, and that they do not provide evidence, let alone a proof, that internal models are necessary. Nathaniel Virgo and colleagues, in a preprint, hold that the argument does not strictly succeed even in its own terms. John Baez has argued that the proof is too easy to be of practical help, and a further critique holds that the notion of model in it is trivial and the proof assumes full state information.[4]

The evidential weight of these differs sharply and should not be flattened. Thobani is peer-reviewed and citable. The rest are a preprint, a blog post and a forum post: legitimate to read, not legitimate to cite as established.

By Chapter 16's standard this lands, and it lands where the tradition has no defense in place. The model leans on the theorem and cannot repair it, since it is not ours to repair, and no amendment has been made. The cost is specific: wherever VSM teaching says that requisite variety or the good regulator theorem *proves* something about organizations, that claim is now contested at the source, and we have not noticed.

It is also the one piece of good news in this part of the inventory. A tradition whose foundational borrowings are being examined by analytic philosophers has, for the first time in years, an interlocutor doing rigorous work on something we care about.

The silence

The adversary that costs most is not a person.

Mainstream organization science does not reject the VSM. It engages Ashby, including through Boisot and McKelvey in the *Academy of Management Review*, and barely engages Beer at all. The VSM literature is concentrated in *Kybernetes*, *Systems Research and Behavioral Science*, *Systemic Practice and Action Research* and the *Journal of the Operational Research Society*; a bibliometric review finds limited diffusion in the managerial literature.[5]

[4] Iqbal Thobani, "A Triviality Worry for the Internal Model Principle," *Synthese* 204 (2024): 36 — peer-reviewed; **quotation unverified.** Nathaniel Virgo et al. (2025), arXiv preprint — **preprint; cite as such or not at all.** John Baez (2016), blog post, and a 2021 forum post on the triviality of the model notion — **grey literature; usable as a pointer to an argument, not as an authority.** Francis and Wonham's internal model principle is the precise control-theory statement and should be cited from the original alongside any of this.

[5] Max Boisot and Bill McKelvey in the *Academy of Management Review* (2010) on Ashby; Francesca Iandolo et al., "Drilling Down the Viable System Theories in Business, Management and Accounting: A Bibliometric Review," *Systems Research and Behavioral Science* 38, no. 6 (2021; online 2020): 738–55, https://doi.org/10.1002/sres.2731, a mapping of 153 Scopus-indexed articles (1990–2018). **Unverified:** the abstract does not state the journal concentration or the "limited diffusion" finding attributed to the review here; check both in the full text. The

Hence the absence of any outside pressure to test. There is no gatekeeper. Nobody's funding, license, accreditation or reimbursement depends on whether the model is right, so nobody has ever been required to find out.

The contrast comes from a constructivist community that was in our position and no longer is. German systemic therapy holds a constructivist self-understanding at least as strong as second-order cybernetics, and it now has an evidence base: two professional associations commissioned a review of more than eighty randomized trials, scientific recognition followed in 2008, an IQWiG report appeared in 2017, and the Federal Joint Committee confirmed benefit for adults in 2018, with statutory coverage from 2020 and for children and adolescents in 2024.[6] The community did not change its epistemology to produce that evidence. An institution that controlled payment demanded it.

Organizational consulting on the same ground has only recently begun commissioning evidence at all, and its external critics describe systemic coaching as an immunization strategy against rational critique, a charge that will sound familiar by this point in the book.[7]

---

claim that mainstream organization science "barely engages Beer" is an impression from a non-systematic look and must be replaced before print by a recorded citation search — the same instrument the Chapter 9 receipt buys.

[6] Kirsten von Sydow, Stefan Beher, Rüdiger Retzlaff and Jochen Schweitzer-Rothers, *Die Wirksamkeit der Systemischen Therapie/Familientherapie* (Hogrefe, 2007), a review of more than eighty randomised controlled trials commissioned by the two professional associations, the Systemische Gesellschaft (SG) and the Deutsche Gesellschaft für Systemische Therapie, Beratung und Familientherapie (DGSF), who submitted the expertise to the Wissenschaftlicher Beirat Psychotherapie (WBP) on 11 July 2006. The WBP, which had refused recognition in 1999, recognized systemic therapy in its opinion of 14 December 2008 (*Deutsches Ärzteblatt* 106, no. 5, 30 January 2009: A208–11). IQWiG, *Systemische Therapie bei Erwachsenen als Psychotherapieverfahren*, final report N14-02, 24 May 2017 (IQWiG-Berichte 513). Gemeinsamer Bundesausschuss (G-BA): benefit and medical necessity recognized for adults, decision of 22 November 2018; psychotherapy guideline amended for adults 22 November 2019, in force 24 January 2020 (billable from 1 July 2020, as recorded in the earlier check of 23 September); for children and adolescents decided 18 January 2024, in force 12 April 2024 (https://www.g-ba.de/beschluesse/3588/, /4028/, /6420/, checked 29 September 2026). The G-BA noted on 12 April 2024 that the service was not yet available in care; state chambers and associations report it billable from 1 July 2024 (not checked in a primary source).

[7] Thomas Bachmann and Johannes Loermann, *Zur Wirksamkeit systemischer Interventionen im Kontext von Arbeit und Organisationen: Ein systematischer Überblick über den Stand der Forschung* (Carl-Auer, 2023), a review commissioned by SG and DGSF — **conclusions not read.** The phrase "immunization strategy against rational critique" is Viktor Lau's (*Schwarzbuch Personalentwicklung*, Steinbeis-Edition, 2013, 29–30), as reported by Waldemar Pelz, "Systemisches Coaching – Kritik am systemischen Ansatz," long web version, https://homepages.thm.de/~hg10086/systemisches-coaching/index.html, accessed 29 September 2026: Lau sees in the use of questionable terms without theoretical grounding "eine systematische Immunisierungsstrategie gegen rationale Kritik" (Lau's book not checked). Short version: Pelz, figshare, 2025, https://doi.org/10.6084/m9.figshare.30119611. Pelz cites Oswald Neuberger, "Ach wie gut, dass niemand weiß, was man so systemisch heißt. Oder die wahre Geschichte des Proteus?," in *Perspektiven systemischer Entwicklung und Beratung von Organisationen*, ed. Nino Tomaschek (Carl-Auer, 2007), for the lack of a definition of "systemic," not for the immunization charge; page range not found.

Beer had seen part of this. His proposal for a standing team of "organizational pathologists," to perform post-mortems on organizations "viable no more" (Chapter 17), describes the work an evidence register would collect.[8] It does not say who would require the post-mortems or pay for them, and that is the piece still missing.

The inference is ours and should be marked as ours: an evidence register without an institutional stake behind it will stay empty. This part hands that finding to Chapter 30, and it changes what the bill is for. The expensive item is not the experiments. It is arranging for somebody to require them.

What this chapter does not establish

It does not establish that the VSM has no serious critics. It has three, and they are in Chapter 16.

Nor does it establish that Snowden is wrong about organizations, or that the complexity critique can be dismissed. It establishes that no test currently separates the two positions, an embarrassment both sides share.

It does not establish that mainstream organization science is right to ignore this model. Its silence delivers no verdict. It shows an absence of attention, and absences of attention are not evidence about anything except attention.

What it establishes is that this tradition has been shadow-boxing. The opposition it imagines (hostile, numerous, ideological) is thin and mostly aimed elsewhere. What has actually kept the model unexamined is that nothing outside it has ever needed the answer.

Receipt — buy a gatekeeper. Approach three bodies that could make a declared failure condition a requirement rather than a courtesy, and secure one commitment in writing: (a) an indexed journal — *Systems Research and Behavioral Science* or *Kybernetes* — to adopt a submission rule that any article reporting a VSM application must state, in a structured abstract field, what result would have counted against the model in that case; (b) a certifying body — Metaphorum, or SCiO, whose Competency Framework route already assesses candidates through a two-hour interview structured as Situation, Objectives, Actions, Results and Aftermath/Learning — to require that the Aftermath/Learning component include one engagement that did not work and what the practitioner concluded from it, and that a submitted diagnosis name its own disconfirming observation. The hook exists; nothing currently hangs on it. Read as published, SCiO's accreditation criteria are a list of positive capabilities — able to use, can use, has developed — with no requirement anywhere to report an unsuccessful engagement, state a method's limits, or say what would have counted against a diagnosis;[9] (c) one public-sector or foundation funder in the DACH region that commissions organisational design work, to require the same in its

[8] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150.

[9] SCiO, "Professional Accreditation," https://www.systemspractice.org/professional-accreditation, **read 14 September 2026**. The SOARA structure, the four grades and the quoted

tender conditions. Price: three conversations and one drafted rule per body, perhaps ten days of someone's time in total, most of it spent finding the right person. This is the cheapest item in the entire programme and the one that determines whether any of the others produce data. If all three decline, the register proposed in Chapter 30 should be built anyway and should be expected to stay empty, and the book should say so rather than pretend otherwise: a programme whose findings nobody is obliged to read is a programme that will be read by nobody.

Handover

This part has said who was standing around the ring. Part Five says what to do in it, and Chapter 30 now has one more line on the bill.

capability phrasing are from that page. **Not checked:** the Competency Framework self-assessment portfolio and CF resources, which are linked as separate documents and are the actual criteria; the claim that no negative-result requirement exists is made here about the published page only, and must be checked against those documents before print. If one of them does require it, the receipt becomes a request to enforce an existing rule rather than to add one, which is cheaper and should be said.

Part IV

What we lose

> **The Ring Without Her** She stops for a while. Between rounds it is fair to ask whether any of this was worth the damage, and whether the club would miss her at all if she did not come back. She thinks it would: watch the fights she is not in, and something is missing that nobody else is attempting. The claim she is left with is narrower than she once thought, and this part states it exactly, and tests whether she is right.

Part Two conceded almost everything. Others found the components of the model independently, described them more carefully, and in several cases measured them. So what remains that is ours? The answer fits in a sentence. Defending it takes five chapters.

Chapter 23, *The Narrow Claim*, fixes the proposition and names four things that do not follow from it. Chapter 24, *Six Things Nothing Else Does*, tests whether the assembly, as distinct from its parts, is unique. Chapter 25, *The Transducer*, places regulatory loss at the crossing of a boundary; the usual remedies, more reporting and more bandwidth, look for it in the channel. Chapter 26, *Freedom as a Quantity*, returns to freedom as Beer stated it in *The Heart of Enterprise*, as something commensurable with the diagnosis, where the tradition has usually treated it as a value asserted alongside it. Chapter 27, *What Might Replace the VSM*, hands each job the model does to whoever does it better and states the residue that has no other home.

One rule governs this part, and we state it as a rule so that nobody has to infer it: every credit carries its concession in the same breath. Nothing is claimed here without what it costs, what it has not been tested against, or what it cannot do, and that goes in the same section, not in a later chapter or a footnote. The reason is not modesty. In our experience this tradition has for fifty years produced lists of virtues with the concessions left off, and Chapter 9 shows that a statement of that form cannot be wrong. To repeat the form here, in a part whose whole subject is what we are entitled to claim, would undo the book.

The stake is real. Part Four is written so that it could come out empty. If, after Chapter 27, the residue is nothing that another literature does not do better, the conclusion follows without appeal: the tradition should be wound up and its useful pieces handed on under other names. We would rather establish that than keep the tradition alive by lowering the test until it passes. A claim narrowed to the point of survival is worth something only if the narrowing was done in public and could have gone the other way.

Chapter 23

The Narrow Claim

The Viable System Model is the only available instrument that will take a named system at a named level and say, in a vocabulary that survives the journey between a person and a planet, what that system needs in order to remain viable — although we have not operationalized that necessity claim, measured it across deployments, or designed the comparative test that could falsify it.

That is the narrow claim, and everything in Part Four is measured against it. The sentence is built so that the claim and the concession cannot be separated: a reader who quotes the first half without the second has not quoted us. Part One ended by conceding that nothing in the model's core, the necessity of the five functions and the recursion claim, has yet been stated in a form that could be disconfirmed. We have now reached the part of the book where we say what the model is good for, and the concession comes with us.

What the claim says

Three phrases carry the weight.

A named system at a named level. Both names matter, and each fails in its own way when dropped. A named system is a particular one — this hospital, this ministry, this cooperative, this fishery — and neither an organization in general nor a class of organizations. The model is applied to a case; it makes no claim to be a theory of firms. A named level is a particular recursion: the ward, or the hospital, or the trust, or the regional system, never all of them at once.[1] Most of the diagnostic disagreements we have witnessed in practice are level confusions. Two analysts agree entirely about the facts and disagree about which system-in-focus the facts belong to. When either name is missing, the model still generates sentences, but nobody can check them, because they have no fixed referent. Demanding both names can look like pedantry. It is the minimum condition under which anything else in the claim can be assessed.

[1] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 116–19. The Recursive System Theorem is at 118: "In a recursive organizational structure, any viable system contains, and is contained in, a viable system." The test for a level follows on 118–19: a division that "could in principle be hived-off as a separate business" is a viable system; the accounting system that prepares the group accounts is not, since it "could not conceivably operate independently of the corporation."

A vocabulary that survives the journey between a person and a planet. This is a claim about the invariance of terms. We are not asserting that a person and a planetary system are alike, or that what is true of one transfers to the other. We assert something weaker and more specific. The same small set of terms (the operational units, the coordinating function, the management of the here-and-now, the intelligence directed outward and forward, the identity-setting function, the channels between them) can be applied at radically different scales without a new vocabulary at each scale. That property is bought with abstraction, and abstraction is exactly what makes a term hard to operationalize. This is why our sentence carries its concessive clause. The feature that lets the vocabulary travel is the feature that makes it hard to hold to account, and we owe the reader that sentence every time we praise the travel.

What that system needs in order to remain viable. This is a necessity claim, and it excludes two things. It is not a prediction: it does not say when the system will fail, or with what probability, or under what load. It is not a recommendation: it does not say that the system ought to continue, or that the arrangement it needs is desirable, or that anyone should build it.

The word is Beer's, and so is its definition. Viable means "able to maintain a separate existence," and Beer refused to let existence stand in for it.[2] Asked whether an organization that exists is "by definition a viable system," he answered: "The fact that it is there does not prove that it is effectively there." The institution whose System Four is empty he called a decerebrate cat, which "lives on" and is sustained "by artificial sustenance."[3] A necessity claim in his sense says only that if this named system is viable at this named level, certain functions are being performed somewhere, by someone or something, however informally and however badly. That last clause is Beer's too. System Two "is always present, or the organization would shake itself to pieces," and "System Four exists — it has to: but it is disseminated."[4] A missing or degraded function then becomes a diagnosable condition rather than a stylistic preference.

The same two moves make the claim hard to embarrass. A persisting organization is read as performing its functions informally, and one that performs them badly is read as sick, not as a system that survived without them. We have not, however, stated the necessity in a form that would let a counterexample count against it; the coding in Chapter 28 is our attempt. A necessity claim that no observation could embarrass is a definition wearing the clothes of a finding,[5] and Part One's closing charge applies to this sentence as much as to any other in the tradition.

---

[2] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 1, quoting the *Oxford English Dictionary*; also Beer, *Heart of Enterprise*, 113.

[3] Stafford Beer, *Think Before You Think: Social Complexity and Knowledge of Knowing*, ed. David Whittaker (Wavestone Press, 2009), 148; Beer, *Heart of Enterprise*, 265.

[4] Beer, *Diagnosing the System*, 66; Beer, *Heart of Enterprise*, 143.

[5] Chapter 9 sets out the criterion we apply here, and Chapter 18 applies it to the published tests of the model.

Why the claim had to be made smaller

Six neighboring literatures, working independently and largely without reference to Beer, arrived at levels, at coordination, at limits on the information a controller can handle, at the autonomy of subordinate units, at monitoring, and at recurrent patterns of persistence. They got there with their own evidence and their own tests, in several cases with more of both than we have. That is Part Two, and it is why the claim is this small.[6]

So we cannot claim that Beer alone discovered the components, and we should not want to. The cost is real, and we would rather state it than have it extracted from us. The discovery claim is gone. With it goes much of what the tradition has said about itself: the singular finding, the laws first seen from one vantage point, the founder's priority that later work merely elaborates. That story has carried a good deal of our community's self-understanding, and it does not survive contact with the neighboring literatures. Chapter 10, *Six Components*, does the accounting component by component. It is not a comfortable chapter.

The reflex to avoid is treating the concession as an attack to be repelled. On any ordinary reading of evidence, independent arrival by several traditions is good news about the components. It suggests they are features of the problem and not artifacts of one modeler's habits. What the concession damages is a claim about origin, not a claim about validity. The two have been run together for so long that separating them feels like a loss, and the loss is of a story.

We are not the first to narrow a claim about the model so that it could be checked. Michael D. Pfiffner, testing the model on 135 organizations, confined crisis to a slowly developing threat to the existence of the whole organization and set aside disasters, sabotage and force majeure, so that he could measure it.[7] His claim, that viability as a whole bears on whether such a crisis occurs, is not our necessity claim. Chapter 28 sets out the coding a necessity claim needs, and Chapter 29 explains why his items cannot supply it.

What replaces it: an assembly claim

What remains is an assembly claim. The VSM puts the components into a closed anatomy, applies that anatomy to a named system at a named level, and returns a diagnosis rather than a list of resemblances. One

[6]The six components and the literatures that reached them are set out, with sources, in Chapter 10; Chapters 11 to 14 take them in turn.

[7]Michael Dominik Pfiffner, "System Viability of Organizations and the Aetiology of Organizational Crisis: A Quantitative Assessment of Stafford Beer's Viable System Model" (PhD diss., Utrecht University, 2017), https://hdl.handle.net/1874/356772, 27 (the scope is "slowly or gradually developing crises"; crises caused suddenly, "in an act of nature or will beyond control", are excluded) and 221 (crisis as a situation "in which an existential threat menaces an entire organization"; "Existential crises from disasters, sabotage or force majeure are not considered"). Checked 29 September 2026 in the calibre PDF, printed pagination.

component it assembles only in part: the panarchy literature's two cross-scale couplings, *revolt* and *remember*, have no counterpart in the assembly, as Chapters 10 and 11 record.

The whole assembly claim lives in that last distinction. A resemblance says: this organization is like that pattern. It maps features onto a template and reports the degree of fit. It is reversible and nearly costless. The mapping can be redescribed, the template swapped, and nothing follows that anyone must act on or could be wrong about. A diagnosis has a different structure: this named absence, at this level, is why this named thing is failing. It has a subject (a particular missing or overloaded function), a location (a particular recursion), and a consequence attached to a particular observed failure. Someone who knows the case can dispute each of the three, and that is exactly what a resemblance never offers. The diagnostic form depends on the anatomy being closed. Because the set of functions is fixed and finite, an absence means something. From an open list, nothing could be missing.

Two qualifications belong in the same breath. First, the diagnostic form is not diagnostic accuracy. We claim that the model produces statements that could be wrong in a specified way, and not that those statements have been shown to be right. Whether our diagnoses hold up when a case is followed forward is something the tradition has not systematically checked. Chapter 29's second experiment, ten interventions with predictions reconstructed from documents written before the outcome and the failures kept, is the check we have not run. Second, this chapter does not settle whether the assembly is unique. That is the job of Chapter 24, *Six Things Nothing Else Does*, which sets each candidate instrument against the assembly one property at a time. Here we fix only what follows if the assembly is not unique: the word *only* comes out of our sentence, and what remains is a useful instrument among several, still worth teaching and still worth testing, but no longer the ground of a distinctive claim.

Four things that do not follow

The claim is narrow at four edges, and at each of them the tradition has been tempted to spread.

The model does not predict when a named system will fail. It identifies conditions whose absence is associated with loss of viability. It supplies no hazard rate, no time horizon and no established threshold. Beer hoped for one: of his identity exercise he wrote that "maybe (one day) thresholds will be established, to exceed which will predict collapse."[8] None has been. A predictive claim would need what predictive claims always need: a defined population of systems, states recorded before the outcome, outcomes recorded afterward, and performance compared against a baseline that

[8] Beer, *Heart of Enterprise*, 432.

does not use the model at all. We have not done this. Reliability engineering, epidemiology and the quantitative parts of organizational ecology do it routinely and do it better. People who need a prediction should be sent to them, and we should stop implying that a structural diagnosis has a temporal edge it does not have.

The model does not yet measure how much viability or freedom the system possesses. The tradition speaks in quantities (variety, attenuation and amplification, balance) and produces outputs that are almost entirely qualitative. Chapter 26, *Freedom as a Quantity*, takes this seriously enough to recover the quantity and show what it would take to compute, and concedes in the same place that no deployment on record has measured it. Until one has, "sufficient variety" is a judgment call by the analyst. We should call it that and stop borrowing the authority of the measure-shaped word.

The model cannot legitimate the boundary the analyst or the client chose. Beer made the choice explicit. The method's first step is "to determine precisely what is the organization to be modelled, and to specify its boundaries," and he gave a rule of thumb for noticing a missing level.[9] The model then works inside the boundary chosen, so every diagnosis it produces is conditional on a choice made before the model was applied, usually by whoever commissioned the work. No property of the anatomy makes that choice correct, and no internal coherence in the resulting diagnosis can validate it afterward. Chapter 4, *Channels Without Owners*, sets out the boundary questions in detail. The Coda states what follows from our having no method of our own for answering them.

The model cannot say what survival is for, or whether this is a system whose survival anyone should want. Viability in Beer is relative to purpose as perceived, and the model takes that purpose as given; it supplies no criterion for choosing among purposes. Its normative clauses, "minimum damage to people and to cost" and the Establishment as a pathology, limit how a system survives, not whether it should.[10] An organization that should be dissolved can be made more viable, and the model will help, competently. Nothing in the anatomy supplies a criterion for refusing the work. That is the Coda's subject, and the gap cannot be closed from inside the model.

Each of the four is a limit on our claim. None is a job to be handed quietly back to the model later, and the difference matters because that handing-back is a recognizable failure mode. When a tradition has no way of being wrong, its scope expands to fill whatever a client wants (prediction,

[9] Beer, *Diagnosing the System*, 2; the rule of thumb, 19: up to "seven or eight" embedded viable systems at the next level down, "not many more, I hope, or else you may be missing a whole level of recursion."

[10] Beer, *Heart of Enterprise*, 158, for freedom as "a computable function of systemic purpose as perceived"; Beer, *Think Before You Think*, 24: "organizations do not have purposes. It is the perception and the purpose of the people who sponsor, operate, use and serve the organization that determine how it shall be deemed effective." The First Principle of Organization asks that varieties "tend to equate" and be designed to do so "with minimum damage to people and to cost" (*Heart of Enterprise*, 565, restating 97); the Establishment as pathological autopoiesis, 412. Checked 29 September 2026 in the calibre PDFs.

measurement, legitimation, ethical warrant), because no test at any boundary stops it. Chapter 9 traces that pattern in our own literature. Naming the four limits here is our attempt not to repeat it, and only an attempt: the discipline holds only as long as we keep restating the limits in the places where clients would prefer to hear otherwise.

The proposition, narrow enough to risk

That leaves a proposition narrow enough to risk: given a defensible boundary and level, the VSM claims a necessary assembly for viability and a vocabulary in which absences in that assembly can be diagnosed. If another instrument already does all of that, or if the assembly cannot be made answerable to evidence, that is, if no coding and outcome of the kind set out in Chapter 28 can be registered for it, our use of *only* fails.

We state our own failure conditions here, in the middle of the book, instead of mounting a defense, and the reason is structural. Part Five prices a set of experiments. An experiment is worth its price only if the claim it tests could lose; otherwise the money buys illustration. A claim that survives every result is merely on display, and the tradition already owns plenty of displays. So the two conditions above are commitments, made before anyone applied pressure: we say in advance what would take the claim from us. Chapter 28, *What Would Count Against It?*, turns them from a statement of intent into specified tests, naming the observations that would settle each one.

The first condition comes due soonest, because it can be assessed against literature that already exists rather than against studies nobody has run. It falls to Chapter 24.

Chapter 24

Six Things Nothing Else Does — and What Each Costs

What Part Two leaves us owing

Chapter 14 set our case record beside Ostrom's, and we did not come out ahead. It was the last of Part Two's comparisons. Chapter 10 had listed the components, Chapter 11 followed nesting through Simon, Koestler and Jaques, and Chapter 12 followed variety logic through Galbraith. Six neighboring literatures, set beside the model, had arrived at most of its anatomy independently. We cannot answer that survey with six claims of discovery. Nearly every part we hold was found by someone else, working from other premises, in several cases earlier and with a better evidential record than ours.

What is left to us is a claim about assembly. The claim to invention did not survive Part Two, and Chapter 23 has already narrowed what remains as far as we think it will go. Here are its six parts in full. Each is a credit we think the model has earned, stated together with what it costs, what it has not been tested against, or what it cannot do. That pairing is the convention of this part, and the reason for it is practical: a credit released without its concession travels further than the evidence behind it, and we have watched that happen to every one of these six.

Each part also carries a maintenance cost, and we mean three specific things by the phrase. Practitioners must learn a vocabulary before the framework does any work at all. The framework makes a narrower set of questions easy to ask, which is the other face of its precision. And a small community must keep spending attention to stop the distinctions from eroding into loose usage. These six cost claims are argued from our own practice and are not grounded in documented cases.

One: a closed set, asserted as necessary

The first is that Beer makes a **closed-set necessity claim**. Galbraith, Mintzberg and Jaques describe what organizations have;[1] Beer asserts what

[1]Henry Mintzberg, *The Structuring of Organizations: A Synthesis of the Research* (Prentice-Hall, 1979); Jay R. Galbraith, *Designing Complex Organizations* (Addison-Wesley, 1973), and "Organization Design: An Information Processing View," *Interfaces* 4, no. 3 (1974): 28–36; Elliott Jaques, *Measurement of Responsibility* (Tavistock Publications, 1956), and *Requisite Organization*

any viable system must have, and puts a number on it. Five kinds of work, no fewer, or the thing stops holding itself together. He argued the closure in logic: "There could not be a 'System Six' — because the five-fold system is a closed system, in logic." He was less sure of the count when he first set it out. Introducing the five systems in *Brain of the Firm* as a hierarchy modelled on the somatic nervous system, he wrote that it "seems to be" five, "because the choice of five (rather than some other number) is somewhat arbitrary."[2]

The shape of that claim matters more than its content. A descriptive typology absorbs a new case by growing a category. A necessity claim absorbs nothing: a sixth indispensable function would end the model rather than extend it. That makes the model refusable as well as serviceable, and refusability is the quality Part Five is written to exercise.

The concession is that a claim this strong ought to be testable and has barely been tested. We have not specified the observation that would show a viable system running without one of the five, and we have not agreed among ourselves what such an observation would look like. Beer's texts make it hard to specify, because he holds that a persisting organization always performs the functions in some form; Chapter 28 sets out the coding that would let such an observation count. Chapter 9 is our account of why that absence never cost us anything. The comparative claim needs the same caution. We say that almost nothing else in organizational thought makes an invariance claim of this shape, but that rests on our reading and not on a survey we have conducted, and we offer it as an impression.

The maintenance cost goes beyond remembering five labels. The community has to keep necessity distinct from a useful descriptive checklist, which is the form the model most often takes in teaching. And a closed set narrows attention to failures that can be expressed as a missing or pathological function. An organization can fail from a cause with no home among the five, and a fluent practitioner will still find a name for it inside the model.

Two: recursion with a criterion

The second is **recursion↗**[3] **with a criterion**. Levels are common property, as Chapter 11 shows at length. What the others lack is the claim that the same anatomy appears at every level, joined to a test for whether a candidate level qualifies at all. In our paraphrase of Beer, the candidate must

(Cason Hall, 1989), as in Chapters 10 to 12. For the partial attestations of Three and Four, see Chapter 14: Ostrom's diagnostic approach for Three, and for Four her design principles, collected from cases and tested in review (Cox, Arnold and Villamayor-Tomás 2010). The contrast is between describing configurations that organizations are found to take and asserting a closed set that any persistent system must have.

[2]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 130; Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 98.

[3]Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

be viable in its own right and one of the units that do what the containing system does.[4] Koestler had the intuition and supplied no test,[5] and the test is what separates a criterion from a resemblance.

The criterion earns its keep by embarrassing work done in the model's name. The ascending order in our companion book *Wild Futures* is, at most of its joints, a ladder of scale and not of recursion, and it was Beer's own test that showed it. A model able to correct a book written in its name is doing more work than a metaphor does.

The concession is that the criterion, applied honestly, excludes a great deal of what people want to do with it, including work of ours that reads better before the test than after. Nor have we applied it systematically across our own case record; the *Wild Futures* result is one application, not a survey. Beer did not always apply it either. His recursions "can be extended upward to the terrestrial globe (within the Universe) and downward to the cell,"[6] a chain his own test would have to check link by link. The maintenance cost is the repeated boundary work of telling a viable recursion from a convenient level, and the opportunity cost of refusing attractive cross-scale comparisons that fail the test. Refusal is unrewarded work, which is why the criterion tends to be quoted more often than applied.

Three: a diagnosis, not a description

The third is that the model **yields a diagnosis rather than a description**. Given a named system, it says which function is missing or pathological and what will follow from the absence. Soft systems methodology yields a learning process; Cynefin yields a sense-making device; Luhmann yields an account; actor-network theory yields a description of unusual richness.[7] None of them walks into a particular firm on a particular Tuesday and

[4]The wording is our paraphrase of Beer and not a quotation. His nearest statements: the operational elements of a viable system exhaust its "basic activities, namely those which exist to do what the system does" (Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 116); a division that "could in principle be hived-off as a separate business" is a viable system, and the accounting system that prepares the group accounts is not (ibid., 118); and "the viable system is necessarily a producer of the organization, and not just an adjunct to it, however important" (Beer, *Diagnosing the System*, 8).

[5]Arthur Koestler, *The Ghost in the Machine* (Hutchinson, 1967); Chapter 11 sets out what his holon supplies and what it does not.

[6]Beer, *Diagnosing the System*, 4. Checked 29 September 2026 in the calibre PDF.

[7]Peter Checkland, *Systems Thinking, Systems Practice* (John Wiley & Sons, 1981), and Peter Checkland and Jim Scholes, *Soft Systems Methodology in Action* (John Wiley & Sons, 1990); Cynthia F. Kurtz and David J. Snowden, "The New Dynamics of Strategy: Sense-Making in a Complex and Complicated World," *IBM Systems Journal* 42, no. 3 (2003): 462–83, and David J. Snowden and Mary E. Boone, "A Leader's Framework for Decision Making," *Harvard Business Review* 85, no. 11 (2007): 68–76; Niklas Luhmann, *Social Systems*, trans. John Bednarz Jr. with Dirk Baecker (Stanford University Press, 1995); Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford University Press, 2005), 137: "If a description remains in need of an explanation, it means that it is a bad description." The one-word summaries are our characterizations, except Cynefin's, which is Kurtz and Snowden's own term (462).

comes back with: your System 2↗ is done by nobody designated to do it, too slowly for these two units, and they will oscillate within the quarter.

That specificity is what exposes the model to correction. A framework that only reframes cannot be caught being wrong. One that names a function and dates a consequence can be, and the practitioner who offers such a diagnosis takes a risk the reframing traditions never have to take.

The concession is that the diagnosis is only as good as a naming the model cannot supply: which system, bounded how, in whose interest. Chapter 5 makes the point for measurement and the Coda makes it for boundaries. Change the naming and the same procedure returns a different diagnosis, with no internal signal that anything has moved. The example above also shows where the model protects itself. For a persisting organization it cannot return "absent" for System Two, which Beer says "is always present, or the organization would shake itself to pieces."[8] It returns degraded presence, unrecognized, slow or overloaded, and degraded presence is harder to show wrong than an absence. The coding in Chapter 28 is meant to tell the two apart. The maintenance cost is disciplined naming before diagnosis, and sustained attention to the evidence that would separate a prediction generated in advance from a label applied after the event. Chapter 15 is candid about how little of that evidence we hold.

Four: a pathology catalog that is generated

The fourth is a **pathology catalog that is generated rather than collected**. System 3–4 imbalance, the autonomy–cohesion oscillation, the algedonic↗ channel that is present but unheeded, a metasystem that has collapsed into System Three: these are derived from the structure and, in Beer's presentation, confirmed in practice.[9] Most rival frameworks we know have a case library where this has a theory, and a theory predicts failures nobody has recorded yet, which a library cannot.

The concession is that generativity without testing is also the recipe for a vocabulary that cannot be wrong. A catalog derived from structure always has a name ready for whatever happened, and the fluency with which we supply the name is no evidence for the derivation. Maintaining the catalog therefore costs prospective coding and failed predictions, where we have so far offered ever more ingenious retrospective diagnoses: the pathology written down before the quarter, the prediction that did not come true recorded next to the ones that did. Chapters 28 and 29 set out what that would take.[10] The one quantitative instrument built on the model has the same gap. Michael D. Pfiffner's viability construct (Chapter 18) was never

[8] Beer, *Diagnosing the System*, 66.

[9] Beer, *Heart of Enterprise*, 175–85, 265 and 408. The collapse into System Three is at 265, reported as observation: "Everywhere one turns, the reality is the same." Checked 29 September 2026 in the calibre PDF.

[10] No published prospective test of a VSM pathology prediction was found (Chapters 15, 18 and 29; searches September 2026). Michael Dominik Pfiffner, "System Viability of Organizations and the Aetiology of Organizational Crisis: A Quantitative Assessment of Stafford Beer's

tested against a baseline or a rival predictor of crisis; that comparison could be run on his 135-case dataset if it is released (Chapter 29). Without it, the framework narrows every failure into one of its own names and reports the narrowing as a confirmation.

Five: one currency across scales and substrates

The fifth is **one currency across scales and substrates**. Variety lets a person, a firm, an economy, a species and a machine collective be discussed in a single vocabulary, and that property is what allows this book to compare claims made at unlike scales at all. Ostrom does not scale down to a person. Active inference does not, at present, scale up to a polity in any usable form. Luhmann does not descend below the social.[11] The convertibility is not a stylistic convenience; it is the condition of the argument we are making.

The concession is that a currency which converts everything may be a metaphor wearing the clothes of a measure. Beer saw it: "Could we ever say: *this* is the variety of the chemical industry, say? Of course we could not."[12] Chapter 3 shows how far the law we quote has drifted from the law Ashby wrote, and Chapter 5 shows what happens when the ladder is used without a ruler. The maintenance cost follows directly: every comparison needs an explicit statement of what counts as a distinguishable state on both sides of it. Where that statement is missing, the shared word hides quantities that do not commute. It also steers attention away from whatever will not convert, which is often the part of the situation that matters most to the people inside it.

Six: the algedonic channel

The sixth is the **algedonic channel**, and it comes close to being singular. The other models we know have no explicit architecture for pain signals that bypass the hierarchy, with a theory of why the bypass must exist rather than merely be tolerated.[13] Control theory, as we read it, has escalation without legitimacy; organizational design mostly has a suggestion box.

The concession is made in Chapter 4, and we do not re-argue it here. Beer says how a signal qualifies, by a statistical filter on points of stability whose criteria the metasystem sets. He did prototype one channel for people outside management, a citizens' algedonic meter on which the

Viable System Model" (PhD diss., Utrecht University, 2017), https://hdl.handle.net/1874/356772, rated crises that had already happened.

[11] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990); for active inference, Chapter 12 and its sources; for Luhmann, *Social Systems* (note above). The judgments about scale are our reading of these literatures, not findings reported in them.

[12] Beer, *Heart of Enterprise*, 86–87.

[13] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 143, 147–49 and 229–30, for the algedonic channel and the argument that a bypass of the hierarchy has to exist rather than merely be permitted.

holder signals anywhere between "total disquiet and total satisfaction" and "does not have to explain anything"; a prototype of ten linked meters was taken to Chile "for experimental purposes."[14] A citizen inside the national recursion is still not an affected non-member. The architecture provides no seat for the affected non-member, and no answer to who chooses the points of stability, who may contest the filter, or who audits it. That chapter was the prosecution. Here the credit comes with the concession attached, and both are true at once.

The maintenance cost is protected attention outside the reporting line. Someone has to specify thresholds, keep the bypass open when the line has every reason to close it, and check whether the people harmed by the system can reach the channel at all. None of that sustains itself, and a channel that exists on a diagram while going unheeded in practice is one of the pathologies the model itself predicts.

The composite, and what it does not settle

The composite is the argument. Each part is attested somewhere else. Nothing else we have found assembles all six. And the assembly is what a book that wants to move between a person and a planet actually requires. That is the whole of the claim, and it is narrower than the claims usually made on the model's behalf.

Uniqueness buys less than it seems to. An assembly can be unique and still wrong. Six credits, each carrying a concession about evidence we have not gathered, add up to a distinctive model and fall short of a validated one. Distinctiveness is a claim about the space of available frameworks; it says nothing about the world. Nothing in this chapter shows that the five functions are necessary, that the recursion criterion is the right criterion, or that a generated pathology outperforms a collected one in a live organization.

Pricing that work is what Part Five exists for, and before we get there the next two chapters take up the transducer and the measurement of freedom. What we are entitled to say here is smaller and, we think, defensible: if the model went, anyone who wanted what the assembly does would have to reassemble these six from six different literatures, and we have not seen anyone attempt it.

[14] Beer, *Brain of the Firm*, 2nd ed., 283–84. Checked 29 September 2026 in the calibre PDF.

Chapter 25

The Transducer

The crossing

A perfectly reliable channel carrying a form with three severity levels will deliver those three levels perfectly, whatever the world on the other side was doing. Ashby counts variety in the channel. Beer adds the crossing. The addition is small in wording and large in consequence, which is why the transducer gets a chapter of its own in Part Four and not a paragraph inside our treatment of the law.

In an organization, information that crosses a boundary has to be translated. Something must take the distinctions available on one side and render them as distinctions the other side can receive, hold and act on. That something is the transducer, and it is rarely a device. In ordinary organizational life it is a weekly meeting at which shift supervisors report to a plant manager; an incident form with a fixed set of fields and a severity dropdown; a sensor whose output is a number where the condition it measures was a texture; a liaison role held by a person who sits in two rooms and carries what she thinks matters between them; a coding scheme that sorts complaints into categories for a dashboard; a summary written at one level for reading at another. Each takes a rich stream of states and returns a poorer one, by rules that are usually implicit and almost never audited.

Beer draws the consequence himself: "channel capacity and transduction capacity" are "not at all the same thing."[1] A channel can carry every signal faithfully and on time while the translation at the crossing has already erased the distinction that mattered. The severity form loses nothing in the channel, and no improvement to the channel will get back what it loses at the crossing.

The claim and its arithmetic

Beer sets this requirement among his four Principles of Organization. The specific claim is that a transducer needs at least as much variety as the

[1]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 47: "The point is to draw a clear distinction between channel capacity and transduction capacity: they are not at all the same thing." Checked 29 September 2026 in the calibre PDF. The confusion is older in his work: "In the channel, or in the transducer?" (*Heart of Enterprise*, 104) and "people rarely distinguish the transduction from the transmission" (129), both checked the same day in the calibre PDF.

channel whose information it translates.[2] Beer's own wording, checked against the first edition, is given in the note.

The arithmetic is simple and, once stated, hard to unsee. If the channel distinguishes ten conditions and the translation collapses them into three categories, seven distinctions are gone at the moment of crossing. More bandwidth downstream cannot restore them. Nor can more frequent transmission, a larger sample, a longer report or a faster refresh. All of these raise the rate and volume at which the surviving three categories arrive, and none recovers the seven. Whatever is lost in transduction is lost before the downstream system has anything to work with, and no downstream investment reaches a loss that happened upstream of it.

Here the principle has teeth. If regulatory failure often sits at the crossing rather than in the channel, the ordinary remedy for a regulator that seems not to know what is happening (more reporting, more bandwidth, more frequent updates, one more dashboard) can be exactly the wrong intervention. It costs money and attention, it produces a visible response to the failure, and it leaves the site of the loss untouched. Worse, it can make the loss harder to see: the organization now has more of the impoverished signal and can point to the volume as evidence of attention. The principle recommends a different first move. Before adding capacity, ask whether the distinctions wanted downstream still exist by the time they get there.

The concessions are substantial, and they belong beside the claim. The observation is Beer's, and we do not present it as ours. The principle has never, to our knowledge, been operationalized or measured. We have no procedure for counting the variety of a transducer, no instrument that returns a number, no threshold, and no published case in which the condition was measured. Beer applied it himself, in the chapter on crisis in *Brain of the Firm*: a party to a conflict loses variety "in transduction into the sensorium — because the transducer may recognize only that variety which the conceptual model itself can discriminate," so that "the variety of its own model decides the variety of the crisis."[3] That is an explanation, not a count. What we credit here is a conceptual identification of a failure site, with no validated measure behind it. That is a real thing to have, and much smaller than the confident language we have sometimes wrapped around it.

[2] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), ch. 4 ("The Exit"), p. 101, the Third Principle of Organization: "Wherever the information carried on a channel capable of distinguishing a given variety crosses a boundary, it undergoes transduction; and the variety of the transducer must be at least equivalent to the variety of the channel." Wording and page checked 23 September 2026 against the calibre copy of the 1979 Wiley first edition. In *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 376–77, Beer restates the requirement in two-sided form, "(iii) Transducers neither attenuate nor amplify variety," as one of four requirements of stability "drawn respectively from the four Principles of Organization" of *The Heart of Enterprise*.

[3] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 364.

Why this matters for the partition problem

Chapter 3 found the same partition problem in three places: who counts the states, who draws the boundary, and who defines the transducer. The third comes due here. The fields on a form, the categories in a coding scheme, the resolution of a sensor and the standing agenda of a meeting each decide in advance which differences may make a difference, and someone with interests makes that decision.[4] Chapter 4's algedonic channel inherits it: a right to raise an alarm through a form with three severity levels is a right to raise one of three alarms. Beer made that channel poor on purpose. The algedonic loop works "on a non-analytic, non-articulated scale of 'pleasure and pain'," and the algedonic controller is "the biggest variety reducer of them all."[5] What has to survive that crossing is the difference between something hurting and all being well, not the detail. Whose distinctions survive a crossing is part of the design, and the Coda returns to it as the book's closing question.

What the claim does not cover

That question also marks the edge of the claim. The nearer candidates should be named one by one, not gestured at as a literature.

Shannon's communication vocabulary has a technical transducer, and it is a transducer in a precise sense that ours is not: defined, measurable, and surrounded by a theory that is quantitative in a way our version has never been. Shannon also states a condition on it for signals. The output of a transducer carries no more entropy than its input, and the same entropy only when the transducer is non-singular, so that the input can be recovered; and the transducer that encodes "should match the source to the channel in a statistical sense."[6] That is close to Beer's condition, stated for signals. What no evidence assembled for this chapter yet establishes is that the organizational requisite-variety condition (that the translating element must match the variety of what it translates) is stated anywhere for organizations, where the transducer is a form, a meeting or a person. The gap is in our searching; we have not demonstrated an absence.

---

[4]Beer makes the classification relative himself. "Many people suppose (without much thought) that because market research 'finds things out' it must be an amplifier of managerial variety"; in his example it reduces the expression of demand to a range the works can handle (Beer, *Diagnosing the System*, 25). Which states count is settled by "whether the resulting change of state serves, or has no bearing on, the purposes of the system" (98). Checked 29 September 2026 in the calibre PDF.

[5]Beer, *Brain of the Firm*, 2nd ed., 283 and 233. Checked 29 September 2026 in the calibre PDF.

[6]C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* 27, no. 3 (1948): 379–423, https://doi.org/10.1002/j.1538-7305.1948.tb01338.x: the discrete transducer is defined in §8, 399; Theorem 7, 400 ("with entropy (per unit time) less than or equal to that of the input. If the transducer is non-singular they are equal"); the matching sentence, 403. Checked 29 September 2026 in the scan of the journal issue (Internet Archive, bstj27-3-379).

Boundary-spanning roles in organization theory pick out the same people and offices Beer's transducer would, and study them with an empirical apparatus we do not have: role studies, measured effects, accumulated cases.[7] We have not established whether that literature attaches a variety condition to the role, or treats spanning as a function to be staffed and supported rather than a capacity to be matched against a load.

Monitoring in Ostrom's institutional work sits at exactly the crossing we are describing and comes with the field evidence Chapter 14 credited. Monitors translate local conditions into terms a governing arrangement can act on, and their adequacy is treated as consequential. Whether that adequacy is analyzed as a variety condition, or as a matter of incentives, proximity and trust, is a question our reading has not settled.

Boundary states in active inference give a formal account of what passes between a system and its environment, with the mathematics we lack.[8] Whether that formalism, applied to organizations and not to the systems it was built for, yields the same requirement is exactly the comparison we have not run.

Each of these is a nearer candidate than the flat claim that nobody else noticed translation, and each stays a candidate, neither refutation nor confirmation, until the comparison is searched and dated. Until then our defensible claim is narrower than the chapter's title might suggest: the Part Two record does not yet contain another framework that joins organizational boundary crossing to a requisite-variety condition. The claim is scoped to the Part Two record as it stands in 2026, not to the literature at large, and should be read with its scope attached. We have not run the negative searches that would settle it. A single citation, from any of the four literatures above or from one we have not thought of, would retire it, and we would rather see it retired by a colleague than defended by us.

One independent use of the word, with a price attached

A small corroboration arrived after this chapter was drafted, from someone with no stake in the claim.

Pfiffner, building the questionnaire discussed in Chapter 18, describes the instrument he made in exactly these terms. The questionnaire, he writes, "in cybernetic terms is a transducer — i.e. it leads from cybernetics to plain business language." And he states what the crossing costs: "In so

[7]For the information-transfer role of boundary spanners, Michael L. Tushman and Thomas J. Scanlan, "Boundary Spanning Individuals: Their Role in Information Transfer and Their Antecedents," *Academy of Management Journal* 24, no. 2 (1981): 289–305, https://doi.org/10.2307/255842. Bibliographic details checked in Crossref, 29 September 2026; we have not read the article.

[8]Karl Friston, "A Free Energy Principle for a Particular Physics," arXiv preprint arXiv:1906.10184 (2019), https://doi.org/10.48550/arXiv.1906.10184, as cited in Chapter 12, whose notes also give the critiques of the Markov-blanket construct (Biehl, Pollock and Kanai 2021; Bruineberg and colleagues 2022).

doing, explanation or diagnostic power is exchanged for more practical relevance."[9]

That is the chapter's claim, reached independently by someone who was using the model, not defending it. Pfiffner does not argue that transduction is where regulatory loss occurs. He reports, as a practical matter, that his own crossing from one variety domain to another cost him something, and he names what. The questionnaire transmitted its content perfectly well; the loss was at the boundary.

It is one instance and not evidence for the general claim, which still needs the coded corpus below. What it shows is the concept doing work for a person who had a boundary to cross and needed a word for what it cost him.

The cheapest test

The test is correspondingly modest. It is small enough to sit beside the claim it tests, so we state it here and not in Part Five.

Receipt — transduction as a locus of regulatory failure. Take a published corpus of documented regulatory failures. Define in advance what counts as lost channel capacity and what counts as loss in transduction, with the definitions fixed before any case is read. Code each case blind to the preferred explanation. Code a second failure type, amplification at the crossing (distinctions added that the source did not make), since Beer's 1981 restatement, given in the note on the Third Principle, is two-sided. Report disagreements between coders as well as counts. Two conditions would cost the claim its force: if failures do not concentrate at crossings, and if coders cannot distinguish crossing from channel without using VSM labels as their evidence. The second is the more dangerous of the two, because a coding scheme that can only find transduction failures by importing the vocabulary that predicts them has measured our habits rather than the world.

A negative result would still be worth having. If failures distribute evenly across channel and crossing, the transducer is one failure site among several, and the priority we give it in teaching is unearned. If the coding cannot be done reliably at all, the distinction between channel loss and transduction loss is not yet operational. That would be a finding about the state of the principle rather than about the cases, and it would tell us that the measurement problem has to be solved before the empirical claim can even be posed. Either outcome beats the present position: a principle asserted for nearly fifty years and never measured.

[9] Michael Dominik Pfiffner, *System viability of organizations and the aetiology of organizational crisis: A Quantitative Assessment of Stafford Beer's Viable System Model* (PhD diss., Universiteit Utrecht, 2017), http://hdl.handle.net/1874/356772, §7.3.1 ("Abstractness of VSM Theory", in the Limitations section of ch. 7), p. 244. **Read in full text 14 September 2026;** page checked 23 September 2026 against the calibre copy of the Utrecht dissertation (2017).

The price is one coding study. We give no figure in money or days, because we have not identified the corpus and cannot honestly price what we have not specified. Identifying and pricing it is the first thing that has to happen if this receipt is to become an experiment and stop being an intention.

Chapter 26

Freedom as a Quantity — The Book He Wrote and We Shelved

The quantity this chapter is about

Chapter 8 records an objection: Cybersyn's architecture distinguished itself from a state bureaucracy by the speed of vertical control rather than by any reduction in hierarchy. Beer answered a version of it himself, and his answer is better than we usually remember and narrower than the objection requires. This chapter settles that debt.

What the answer lacks is a measure, and the measure is the item in Part Four hardest to price. Chapter 27 lists it among the residue: freedom treated as a quantity commensurable with the diagnosis, a magnitude that falls out of the same arithmetic that sizes a channel or places a decision. The concession comes with the credit. We call the quantity *residual variety* in these pages: the variety left to a person or an operational unit after an institution's attenuation has taken what the institution needs to stay stable. The term is Beer's, and he used it for something else. In *The Heart of Enterprise* residual variety is what the balance between Systems Three and Four leaves over for System Five to absorb, the quantity formalized in his Third Axiom of Management. Ours is the room left at the other end of the design, which Beer in the same book called freedom. We keep the term because it names a residue that a design either preserves or destroys, and in this book it carries only that sense.[1] Beer had declared the quantity computable in principle, as the section on *The Heart of Enterprise* below

[1]The phrase *residual variety* does not occur in *Designing Freedom*; the concept there is unnamed, and we have supplied the term. The phrase is Beer's own, but from elsewhere in the corpus and for a different quantity. In *The Heart of Enterprise* (John Wiley & Sons, 1979) the metasystem's variety absorbers are there "to contain the residual variety not absorbed by the first three" (216), and System Five "necessarily absorbs the residual variety of the Three-Four interaction" (263). The usage is formalized in the Third Axiom of Management, 298: "The variety disposed by System Five is equivalent to the residual variety generated by the operation of the Second Axiom." The appendix restates the axiom with "equals" in place of "is equivalent to" (567), as does *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 130 and 134, where residual variety is what is left over from the operation of the Second Axiom for System Five to absorb. Beer's term for the person-and-unit quantity is freedom: "freedom for the operational unit is a computable function of the purposes of the system as perceived" (*Heart of Enterprise*, 200). Part Four uses *residual variety* throughout in the person-and-unit sense of the lectures, as defined in the text, not in the Third Axiom sense, and the distinction should be kept in view wherever both books are cited on the same page. Pages checked 29 September 2026 against the calibre copies of both books.

shows. He gave no procedure for computing it, and no deployment has measured it.

What *Designing Freedom* actually argues

Six radio broadcasts, autumn 1973, the thirteenth series of Massey Lectures for the Canadian Broadcasting Corporation.[2] Beer wrote the first of them in a house on the Chilean coast.[3] The audience was a general one. The argument was not.

The first lecture treats institutions as dynamic systems, whose failures are outputs of their organization, "not accidental; and they are not mistakes."[4] Every institution faces more variety↗[5] than it can absorb, so under Ashby's Law↗ every institution reduces variety: a boss constrains his subordinates, rules constrain interactions, and "someone shoots the cat," forcing those the institution deals with into stereotypes. "This works, but is no fun if you are the cat."[6] What is left to a person after that reduction is the quantity we call *residual variety*; the lectures do not name it. Beer draws the consequence in the same paragraph: "It is in this way that freedom starts to be subordinated to efficiency; but the only alternative — which we must

[2]Stafford Beer, *Designing Freedom*, Massey Lectures, 13th ser. (Canadian Broadcasting Corporation, 1974). The book is the text of six radio broadcasts given in the autumn of 1973 as the thirteenth series of Massey Lectures, established by the Canadian Broadcasting Corporation in 1961. The lectures are: 1, "The Real Threat to 'All We Hold Most Dear' "; 2, "The Disregarded Tools of Modern Man"; 3, "A Liberty Machine in Prototype"; 4, "Science in the Service of Man"; 5, "The Future That Can Be Demanded Now"; 6, "The Free Man in the Cybernetic World". A John Wiley edition (London and New York) carries the same 1974 date, with a further Wiley printing dated January 1975; House of Anansi Press reissued the book in 1993. Earlier notes for this book gave the Canadian publisher as CBC Publications and the Wiley edition as 1975; the imprint appears in catalogue records variously as Canadian Broadcasting Corporation, CBC Publications and CBC Learning Systems, and we have not seen a title page, so we give the corporate name. Quotations in this chapter were verified against the full text of all six broadcasts and are cited here by lecture; the print edition runs to about a hundred pages, but page references for it have never been established in our notes and are still to be supplied. We have not invented any.

[3]Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 197–98. Beer returned to Chile on 20 June 1973 and spent six weeks in the coastal town of Las Cruces, partly for his own safety, drafting the Massey Lectures there. The first lecture opens on the setting in Beer's own words — a little house on a steep hill "in a quiet village on the western coast of Chile."

[4]Beer, *Designing Freedom*, lecture 1, "The Real Threat to 'All We Hold Most Dear' ". The sentence in full: "These are not accidental; and they are not mistakes." Page reference to be supplied from the print edition.

[5]Terms marked ↗ are defined in the Glossary at the back of the book. The mark is set on the first appearance of each term in each chapter, and not afterwards.

[6]Beer, *Designing Freedom*, lecture 1, "The Real Threat to 'All We Hold Most Dear' ". Beer's own list of "three of the main ways by which institutions reduce their variety" gives the third as "Someone shoots the cat: the institution does not accept arbitrary interference, and forces those with whom it interacts into stereotypes"; the aside quoted here is from the body of the same lecture. Page references to be supplied from the print edition.

face — is total anarchy."[7] He is not against variety reduction. He is against variety reduction that nobody chose, examined, or is accountable for.

The second lecture turns to the citizen. No insurance company can provide requisite variety for a unique human being, so "your variety must be attenuated, cut down so that your case fits a more general case." Beer accepts the attenuation and objects to the pretense: institutions "should acknowledge, at least to themselves, that they are satisfying Ashby's Law by variety attenuation."[8] In this book *attenuation*↗ names something done to people, and the residue left after it is what they have to live on.

The third lecture is the constructive one. Government is "a kind of machine meant to operate the country in the interests of individual freedom", and it does not work, so "let us redesign this 'liberty machine' to be, not an entity characterized by more or less constraint, but a dynamic viable system that has liberty as its output."[9] A *liberty machine* is an institution whose output is the residual variety of the people inside it, engineered on purpose. Beer had used the phrase three years earlier, as the title of a 1970 address in Washington, for a different principle: liberty redefined so that "competent information is free to act."[10] Between 1970 and 1973 the term moves from a property of information flow to a property of people.

The answer to the control-room objection, and its limit

Beer then states the objection that has followed the tradition ever since: this is the death of individual freedom, a control room with a dial for citizens. His answer deserves quoting more often than we quote it. First, the machinery of government is not the policy of government, and the argument is about the "how" of governing rather than the extent of intervention. Second, and sharper: "It is perhaps more alarming that private concerns

[7] Beer, *Designing Freedom*, lecture 1, "The Real Threat to 'All We Hold Most Dear' ". Page reference to be supplied from the print edition.

[8] Beer, *Designing Freedom*, lecture 2, "The Disregarded Tools of Modern Man". Both quotations in this paragraph are from that lecture. Page references to be supplied from the print edition.

[9] Beer, *Designing Freedom*, lecture 3, "A Liberty Machine in Prototype". Both quotations in this paragraph are from that lecture. Page references to be supplied from the print edition.

[10] Stafford Beer, "The Liberty Machine", an address to a conference on ecological systems organized by the American Society for Cybernetics, Washington, DC, 8 October 1970, reprinted in *Platform for Change* (John Wiley & Sons, 1975), 307–21; the redefinition of liberty is at 319. We first had this passage and the one quoted later in the chapter at second hand, from Werner Ulrich's citation of them in "A Critique of Pure Cybernetic Reason: The Chilean Experience with Cybernetics," *Journal of Applied Systems Analysis* 8 (1981): 33–59, the paper discussed in Chapter 16. Both were checked on 23 September 2026 against the calibre copy of *Platform for Change* (the Wiley "Stafford Beer Classic Library" reprint of 1994, whose pagination agrees with the pages Ulrich cites from the 1975 edition). Its headnote describes the piece as the keynote address to the Conference on Ecological Systems, the American Society for Cybernetics' annual national conference, delivered in Washington, D.C., on 8 October 1970, and reprinted "by permission of Spartan Books". One detail of Ulrich's quotation is worth recording: where Beer writes "the new Liberty Machine", Ulrich inserts the bracketed gloss "[Cybersyn?]", which is his reading and not Beer's text.

are able to build systems of this type, without anyone's even knowing about their existence, than that democratically elected governments should build them in open view and with legal safeguards."[11] That was 1973.

The reply is worth having, and it does not reach the objection Chapter 8 records. Beer answers a claim about purpose and accountability: who builds the instrument, under whose mandate, visible to whom. The Chapter 8 objection is about architecture, and it survives in a narrower form. Beer did answer it architecturally, in *Brain of the Firm*: each level receives the whole Cybersyn package, its autonomy "preserves it from nosey-parkerdom," and "any given level of recursion receives directly only its own Cyberstride reports."[12] What neither book supplies is a measure that would show whether the rule held, and Beer himself reports that practice was "not always" in the intended spirit. A channel carrying production data upward faster than a ministry could is a faster vertical channel whatever the mandate. The lectures supply no measure that would tell an accelerated hierarchy from a design that returned room to the people inside it. That measure is the quantity this chapter is about.

The fourth lecture draws the fault line explicitly, showing the same tools assembling either way, and the sixth states it flatly: "Science is not a neutral thing, as many scientists themselves try to believe."[13] The technology does not choose. Someone chooses, and the choosing is part of the design, not a stage after it. That leaves the question the tradition has never answered in method: whose choosing it is.

The management book that kept the word

Between the broadcasts and the handbook stands *The Heart of Enterprise* (1979), written for managers. It has a chapter called "Freedom," and it does in the model's own terms what the lectures did in political ones. In the viable system, Beer writes, "the dear question of freedom is locked into its orthogonality," the interaction of horizontal and vertical variety "which

[11] Beer's statement of the four criticisms, and the reply that distinguishes the machinery of government from government policy and confines the argument to the " 'how' of governing, not . . . the extent of intervention", are in lecture 3, "A Liberty Machine in Prototype". The sentence quoted here about private concerns is not: it is the closing sentence of lecture 2, "The Disregarded Tools of Modern Man". We have put the two together because they answer the same objection, but a reader following the argument through the book will find them a lecture apart, and the second is stated before the objection it answers. Page references to be supplied from the print edition.

[12] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 275; see also 288–89 for the Chilean principles, and 271 and 322 for the concessions.

[13] Beer, *Designing Freedom*, lecture 6, "The Free Man in the Cybernetic World"; in full, "So science is not a neutral thing, as many scientists themselves try to believe." Earlier drafts of this chapter placed the sentence in lecture 4, "Science in the Service of Man". Lecture 4 is where the two assemblies are set against each other — "The societary use of science we have is threatening; it becomes oppressive and alienating. The societary use of science we could have is a liberation" — but the flat statement that science is not neutral comes two lectures later, and we have corrected the attribution. Page references to be supplied from the print edition.

DEFINES FREEDOM within the viable system." The reality of freedom "lies, not in the subjective impressions of the people involved, but in what the institution DOES to them by way of variety constraint." A seven-step argument that he calls "basically mathematical" ends in the sentence this chapter turns on: "*Freedom is in principle a computable function of systemic purpose as perceived*." The argument, he adds, can be extended to individual freedom within a society.[14]

So Beer made freedom a quantity in the model himself, in a book addressed to managers. He did not compute it or give anyone a procedure for computing it. The clause "as perceived" is where the question of standing enters his text, since purpose is "a subjective phenomenon, rather than a property of the system independent of its instigators, participants, and observers"; whose perception governs, he does not say.[15] The subtitle of this chapter should be read with that in mind. Beer formalized the quantity and left it uncomputed; the tradition shelved both the lectures and the formal version, and the shelving is ours.

Two claims of standing

In the last lecture Beer says that requisite variety for running the world does not exist in any one head. "Requisite variety for running things properly exists with the people who generate the world's variety in the first place, and that means everyone. Whoever opts out of his or her regulatory role is robbing the total system of its power to be stable. Therefore it is not for me to specify the content of the total regulatory model, but only to point to the need for it. But if this stricture applies to me, it also applies to you. The requisite variety for being messianic belongs only to the genuine Messiah."[16]

The passage does two things at once, and we have generally noticed neither. It refuses to specify what the system should be for, and it grounds the refusal in requisite variety, where one might have expected modesty. In the same breath it makes a strong normative claim: the people affected must settle the content, and abstention, far from being neutral, is theft from the system's stability.

One complication belongs here rather than in a footnote. Standing for those who generate the variety is not the only claim of standing Beer made. In the 1970 address where *liberty machine* first appears, he made a different one, on behalf of the profession: "Who are we to say that we know best and everyone else is wrong? … It is that we are responsible. We are not responsible because we have been elected to govern affairs; we are responsible because cybernetics, that science of effective organization, is

[14] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 145–46, 156, 157–58 and 160; the computability sentence is at 158 and restated at 200.

[15] Beer, *Heart of Enterprise*, 158.

[16] Beer, *Designing Freedom*, lecture 6, "The Free Man in the Cybernetic World". Page reference to be supplied from the print edition.

our profession. Such understanding of this subject as there is, we have. Therefore we must speak out."[17]

Both sentences are his, three years apart, and they do not sit easily together. We kept the one about our profession and let the one about everybody else's standing go quiet. That is our admission, diagnostic authority without an account of who may invoke it, and it arrives from the founder rather than from his critics. Chapter 16 reaches it from the other direction, in Ulrich's hands.

What happened after 1973

Beer did not go quiet. The convenient story has the political phase ending with the coup and the man retreating into method. That story is false, and it would let us off.

The Chilean project ended on 11 September 1973, in the coup that killed the government that had commissioned it.[18] In the last lecture Beer did not say that the technology had failed or that the design was premature. His reading was this: "the rich world would not allow a poor country to use its freedom to design its freedom."[19] Chile appears in this chapter as history and never as evidence. A tradition that has lived off one interrupted experiment for fifty years should not be asked to accept a sixth retelling of it as an argument, and the retelling would not establish the quantity in any case. What matters here is that on that occasion the political question was settled outside the model, by force, which is a boundary problem the Coda returns to.

---

[17] Beer, "The Liberty Machine" (1970), in *Platform for Change*, 321; source as in the note above. The ellipsis marks the omission of two sentences from Beer's text; it is not in the original. Ulrich quotes the passage in 1981 and answers it with C. P. Snow, *Science and Government*, The Godkin Lectures at Harvard University, 1960 (New York: Mentor / New American Library, 1962), 118: "If you are going to have a scientist in a position of isolated power, the only scientist among non-scientists, it is dangerous whoever he is." Snow's point, as Ulrich reads it, is not that the advice is bad. It is that the belief that value-neutral advice in political affairs exists — held by the adviser, and worse, by the non-specialists around him — is itself the danger. Chapter 16 establishes that warning as a condition on diagnostic authority.

[18] For the history of Project Cybersyn and its end we follow Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), rather than the tradition's own recollections; on the coup, 208–9. The biographical details in this section — the 1974 withdrawal, the move to Wales by 1976 — are hers, at 223–26: "In 1974 he embarked on a journey of spiritual and material reinvention, first by taking extended trips to the Welsh countryside. By 1976 he had relocated permanently to a small cottage in Wales that lacked running water." The quotation is at 224. Page numbers checked 23 September 2026 against the calibre copy of the MIT Press print edition (2011). The Cwarel Isaf Institute's brochure *Stafford Beer's Cwarel Isaf Cottage: History and now*, in the section by Allenna Leonard, says that Beer bought the cottage, Cwarel Isaf in Ceredigion, "and moved there in 1975", and that it was his base from then until 1985, when he also bought a house in Toronto. The brochure quotes Beer's own description from his preface to *I Said You Are Gods*: "It is very small, and has no piped water, but it is robust." A move in 1975 is consistent with Medina's "by 1976".

[19] Beer, *Designing Freedom*, lecture 6, "The Free Man in the Cybernetic World". Beer's account of the Chilean project and its destruction occupies the closing pages of that lecture, broadcast within weeks of the coup. Page reference to be supplied from the print edition.

Nor did Beer withdraw. *Platform for Change* appeared in 1975 with its politics intact, though it collects material from 1969–73; the evidence that he kept them after 1973 is the freedom chapter of *The Heart of Enterprise* (1979, ch. 6) and the Chile chapters of *Brain of the Firm* (1981).[20] In 1974 he began the retreat to Wales, and by 1976 he was living in a stone cottage without running water. Two decades later "World in Torment" still traced starvation, war and the exploitation of nature and of indigenous peoples to attempts to overcentralize the control of complex systems.[21]

So the depoliticization is no story of a man losing his nerve. The story is about which of his books the method descends from, and there the finding is uncomfortable, because the answer is his.

The two books

Set *Designing Freedom* (1974) beside *Diagnosing the System for Organizations* (1985). They are almost the same length, roughly 29,700 words against 34,200, and the variety logic is just as dense: 255 occurrences of *variety* in the first, 266 in the second. Nothing has been given up analytically, and the later book adds the axioms and the charts. But *freedom* falls from 113 occurrences to 12, *liberty* from 23 to 1, and *management* and *manager* together rise from 23 to 219.[22] The counts are raw string frequencies from the editions we hold. They treat inflections as stems, and the later column has not been re-verified against a second copy. The argument rests on the direction and size of the divergence, not on the figures.

The apparatus is preserved perfectly. What changes is the book's job, and with it the referent of every term the earlier book had made political. The addressee alone does not explain the change: *The Heart of Enterprise*,

[20] Stafford Beer, *Platform for Change* (John Wiley & Sons, 1975). Its addresses and papers are dated from 1969 to February 1973 (the Richard Goodman Memorial Lecture, 14 February 1973); checked 29 September 2026 in the calibre copy (1994 reprint). The freedom chapter is *The Heart of Enterprise* (John Wiley & Sons, 1979), ch. 6, "Freedom", 145–72; the Chile chapters of *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), are chs. 16–19, 245–348.

[21] Stafford Beer, "World in Torment: A Time Whose Idea Must Come", *Kybernetes* 22, no. 6 (1993): 15–43, https://doi.org/10.1108/eb005989. The typescript circulated by Metaphorum is headed December 1992. Earlier notes for this book dated this address to 1990; the 1993 *Kybernetes* publication is the citable version, and any claim about the occasion on which it was first delivered needs separate verification before print — we tried and could not establish it.

[22] Stafford Beer, *Designing Freedom* (Canadian Broadcasting Corporation, 1974), and Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985). Counts are from machine-readable text of the two editions we hold and are indicative rather than exact: they are raw string frequencies, include front and back matter, and count *politic*, *manage* and their inflections as stems. The figures given are *freedom* 113/12, *liberty* 23/1, *variety* 255/266, *management* 10/146 and *manager* 13/73, on totals of 29,719 and 34,232 words. Recounting the *Designing Freedom* column against an independently typeset full text of the six broadcasts gives *variety* 255, *manager* 13, *management* 12, *freedom* 106 and *liberty* 22 on 29,370 words — the same figures within the noise that running heads and typesetting introduce. We have not been able to re-verify the *Diagnosing the System for Organizations* column against a second copy. The point does not rest on the precise numbers, which anyone can recompute, but on the direction and size of the divergence: an order of magnitude in the political vocabulary and none at all in the technical one.

written for managers six years earlier, kept freedom as a chapter title. The handbook declares its own narrowing. It sets aside psychology, social anthropology and industrial relations, and the problems of measurement and filtration, in order to treat structure as briefly as possible.[23] Freedom survives in the handbook as the freedom of a box: what the First Axiom leaves over "means the freedom remaining to the management on the horizontal axis to manage."[24] Oppression survives too, as a diagnostic phenomenon. In a genuinely acute passage Beer notes that with the axiom satisfied and the objective facts identical, "System Three feels saintly, System One feels oppressed."[25] He offers the observation as a symptom to be read off a structure. The handbook does ask the diagnostician to check the First Axiom "as System Three sees it; as System One sees it," and promises "ALWAYS surprises,"[26] but System One's view enters as data for that check and not as a claim that could change the design. The handbook also keeps *residual variety*, in Beer's Third Axiom sense: what the Three–Four balance leaves for System Five↗ to absorb. That is a different quantity from ours under the same name, and it never referred to a person. For freedom and oppression, the concept is intact and the person has become a source of data.

This was no oversight. The handbook recommends the substitution, repeatedly, as a gain. Do not be "bullied by current practice or power politics" when identifying the systems in focus. Naming a mechanism an anti-oscillatory device "makes the issue of potential litigation open to discussion in cybernetic terms. Otherwise it is an emotive and political matter." The theory of autonomy and decentralization begins in a fact about variety "rather than in political theory." Each time, the political description is the inferior one and the cybernetic reading is offered as an escape from it. The offer has substance: a variety argument travels across a boardroom where a political argument stalls, and that is a real gain in a real room. Its cost shows in the one place the handbook lets the citizen back in, a passage on preventative laws, curfews and identity papers as attenuators↗ of public variety, where the question "How does this bear on freedom?" is set as an exercise for the reader and left unanswered.[27]

[23] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), ii.

[24] Beer, *Diagnosing the System for Organizations*, on the First Axiom of Management, p. 102, where the boxed term being defined is *autonomy*. This page and those in the two notes that follow were checked 23 September 2026 against the calibre copy of the Wiley edition (1985).

[25] Beer, *Diagnosing the System for Organizations*, p. 97.

[26] Beer, *Diagnosing the System for Organizations*, 95; see also 99 on the declared and imputed purpose that "converge on a compromise purpose."

[27] Beer, *Diagnosing the System for Organizations*. The three quotations in this paragraph — on not being "bullied by current practice or power politics" when identifying the systems in focus; on naming a mechanism an anti-oscillatory device and what that does to a question of litigation; and on the theory of autonomy and decentralization beginning in a fact about variety "rather than in political theory" — and the passage on preventative laws, curfews and identity papers as attenuators of public variety, are all from that book, at pp. 13, 80 and 37–38 respectively, and p. 34 for the passage on preventative laws.

We are describing two books, not conducting a survey, and the claim should be read at that weight.[28] The direction of descent, though, is not in serious doubt: the method literature descends from the handbook, the one book that declares it leaves these questions aside, and not from the broadcasts or from *The Heart of Enterprise*, and it inherited the substitution along with the axioms. Our critics noticed before we did. Michael Jackson's placement of organizational cybernetics is the charge: the VSM "is most at home with 'complex-unitary' contexts" and "struggles with 'people complexity'." Chapter 16 treats it as a placement rather than a refutation.[29] The shape of the reply says more. The field's own survey of criticism recognizes two lines of objection, both normative, and treats them as accusations to be rebutted.[30] Here the political question survives mainly as something done to us in review, hardly at all as a research program. None of the three books asks a practitioner whose residual variety an intervention should increase, although *The Heart of Enterprise* poses the dependency on perceived purpose, and a dated search of the method literature is needed before that absence is claimed of the literature as a whole.

What the substitution costs

The cost is practical: a tradition that cannot say what it is for cannot recruit anyone who is not already inside it. Chapter 30 prints the two dated counts (as of August 2026). We do not re-derive them here, and the

[28]The depoliticization claim in this section is supported by a close reading of two primary texts by Beer himself, plus the reception evidence cited in Chapter 16 and the companion research register. It is not a systematic survey of the VSM method literature. Such a survey — with a stated inclusion rule, coding for whether the placement of authority over the model is treated as a choice or as a convenience, and a count — would either strengthen this claim considerably or refute it, and it has not been done. It belongs among the claims at risk in Part Five and in the companion register.

[29]Michael C. Jackson and Paul Keys, "Towards a System of Systems Methodologies", *Journal of the Operational Research Society* 35, no. 6 (1984): 473–86, https://doi.org/10.1057/jors.1984.101, developed in Michael C. Jackson, *Systems Thinking: Creative Holism for Managers* (John Wiley & Sons, 2003) and Michael C. Jackson, *Critical Systems Thinking and the Management of Complexity* (John Wiley & Sons, 2019). The wording quoted here is Jackson's own, from *Critical Systems Thinking and the Management of Complexity*, 334 and 336 (checked 29 September 2026 in the calibre PDF); there he also calls the VSM "unsuitable" for the issues of coercive contexts (337). The phrase in which the charge usually circulates, "unable to deal with individual, social and political issues", we have found only as Angela Espinosa's quotation of Jackson (note below) and not in this book, and we do not use it.

[30]Angela Espinosa, "Revisiting the Viable System Model as an Emancipatory Systems Approach", *Systems Research and Behavioral Science* 42, no. 1 (2025): 171–88, https://doi.org/10.1002/sres.3090, published in a Festschrift for Michael C. Jackson. Espinosa takes up two of Jackson's criticisms — that the model is unitary and functionalist and so "unable to deal with individual, social and political issues", and his placement of Team Syntegrity — and argues that Beer's successors have addressed several of them. Reading this as *the* survey of criticism, and reading its rebuttal posture as evidence that the political question survives in the tradition mainly as an accusation to be answered, is our inference from the shape of the literature and not a finding Espinosa states.

searches have not been re-run, so they should be read as a snapshot.[31] Those numbers are usually treated as an adoption problem, to be fixed with better exposition. Read them again with this chapter in hand. The people now building the systems that will attenuate everyone's variety for the next fifty years have anatomy enough. What they lack is an argument about what the anatomy is for, and they are looking for one. We have had that argument in the drawer since 1973. We have not turned it into a measure, a method step, or a question anyone is required to ask on site, which is why having it has so far cost the field nothing and bought it nothing.

What the quantity does that the neighbors cannot

We keep Beer's treatment, rather than borrowing one, for its mechanism. Beer has variety, attenuation, amplification, channel capacity and, as Chapter 25 established, transduction. He can therefore say *how* a design destroys someone's room: which distinction was dropped at which interface, and where the room went. That is more than the surrounding literatures offer. It still stops short of a criterion, because the mechanism has never measured the harm of a single deployment and cannot say whether the room destroyed was worth having. The capability literature carries the evaluative content and is poor at mechanizing the loss; our variety count can mechanize the loss and cannot evaluate it. Neither is a criterion alone, and the honest form of the credit claimed in Chapter 24 is that we hold one half of a pair.

The same logic separates alerting from reporting. Beer's design problem was never a shortage of alerting channels. The trouble was that the alerting channel is subordinated to the reporting channel, so bad news travels the slow path and arrives pre-attenuated by the people it indicts. That is a statement about whose variety is absorbed by whom, and therefore a statement about residual variety. The formulation is our compression of an argument Beer makes in *The Heart of Enterprise*: management information keeps calm by aggregating, and so suppresses the signals that point to instability, which is why "The alerting system is orthogonal to the calming system."[32]

[31] Both figures are printed, dated August 2026, in Chapter 30, "The state of the count"; the searches that produced them are recorded in the companion research register. They are not re-derived here. We have not re-run the searches for this chapter, and counts of this kind move.

[32] The term is Beer's. See Stafford Beer, *Brain of the Firm: The Managerial Cybernetics of Organization*, 2nd ed. (John Wiley & Sons, 1981), on the algedonic signal and its channel, and *The Heart of Enterprise* (John Wiley & Sons, 1979) on why an alert that travels the reporting line arrives filtered. The formulation used here — that the design fault is the subordination of the alerting channel to the reporting channel rather than any shortage of alerting channels — is our compression of that argument and not a quotation. For the argument, see *The Heart of Enterprise*, 406–7, and *Brain of the Firm*, 178–79, where algedonic information must not be "simply absorbed" by the homeostat, and 275–76, where Beer asks what happens if a unit in trouble is "too cocksure, or too lazy, or too corrupt to mention its problem" and answers that

What would settle it

Three things would turn this chapter's reading into a finding. The first is the survey that the two-book comparison stands in for: a systematic pass over the method literature with a stated inclusion rule, a coding scheme for freedom and management vocabulary and for where authority over the model is placed, and reported counts. It would either strengthen the depoliticization claim or refute it, and Chapter 30 prices it instead of repeating the charge.

The second is a method step. Chapter 23 makes the narrow claim that the model tells you where to look. The addition asked for here is a required question at diagnosis: whose residual variety should the intervention increase, and by what observable? The answer is recorded before the intervention, not after. Nothing in the axioms forbids the question, and nothing in the practice requires it. Beer's own formula points to it: freedom is a function of purpose "as perceived."

The third is a measurement, and it would settle the Chapter 8 objection instead of arguing it. Until someone counts the distinctions available to a unit before and after a deployment, the claim that a viable design returns room to the people inside it, rather than merely accelerating the vertical channel, remains a claim about intent. A VSM quantity can be defined and published; Michael D. Pfiffner did it for viability. His score sums twenty-two rated items to a maximum of 176, and every organization in his sample that scored 92 or lower had been through an existential crisis.[33] It is a quantity of viability, not of freedom, and Chapter 18 shows that some of the items it is tested against sit on the outcome side. We have made it for fifty years on the strength of a book we shelved and an experiment that was ended by force.

the alert converts automatically into an algedonic signal (checked 29 September 2026 in the calibre PDF).

[33] Michael Dominik Pfiffner, "System Viability of Organizations and the Aetiology of Organizational Crisis: A Quantitative Assessment of Stafford Beer's Viable System Model" (PhD diss., Utrecht University, 2017), https://hdl.handle.net/1874/356772, 192 (twenty-two viability questions of eight points each, maximum 176; "all 39 organizations scoring 92 and below … experienced an organizational crisis") and 234. Checked 29 September 2026 in the calibre PDF, printed pagination.

Chapter 27

What Might Replace the VSM

The handovers

For mathematics, Aubin, Powers and Friston. For evidence about multi-level governance, Ostrom. For transformation and phase, panarchy and process theorists such as Tsoukas and Chia, whom Chapter 7 cites. For meaning, Bateson, Vickers and Luhmann, whom Chapter 6 introduces. For asking whose boundary is in force, Ulrich, as Chapter 16, *The Objections That Land, and What They Cost*, established and the Coda, *Judy*, will put into practice.[1] That is what actually replaces this tradition, and a reader who needs one of those jobs done should reach past this shelf. In this book, the success of those handovers is nearly an argument against the tradition.

Nearly, because each handover is one-to-one, and the claim we inherited was not. Each neighboring literature takes a job the Viable System Model also claims and does it with instruments we do not have: a formalism, a comparative dataset, a phase model, a theory of meaning, a procedure for exposing whose interests a boundary serves. A literature does a job better, in the sense used here, when it does it with an instrument someone outside it can check: a formalism that computes, a coded comparative dataset, or a procedure others can run. If the tradition claimed only that these jobs matter, the handovers would settle the matter against us. The jobs would be better housed elsewhere, and the model would be a redundant index of other people's work. The argument stays open because the tradition also claimed something no single handover addresses: that a particular closed set of functions must be present together, that the losses between them can be located, and that one measure travels across the whole. The handovers failing would not establish that claim, and their succeeding does not refute it. It has to be tested on its own, in the rest of this chapter and in Part Five.

Reaching past the shelf has a price, and we are asking readers to pay it. It usually means learning a formalism rather than a vocabulary. Viability theory asks for the mathematics before it gives anything back. Active inference wants a generative model of the system before it will compute. Ostrom's design principles ask for a case corpus and the patience to code it, and Ulrich's boundary questions ask for a room in which the affected parties can actually be brought into the discussion. Each time, the practitioner

[1]Sources for each handover are in the owning chapters: Bateson, Vickers and Luhmann, Chapter 6; Ostrom, Chapter 14; panarchy, Chapters 5 and 11; process theory (Tsoukas and Chia), Chapter 7; Ulrich, Chapter 16. For the mathematics, see the note below.

gives up what the diagram supplies at once, a shared picture in an afternoon, for something the diagram cannot supply at all. For the jobs listed above we think the trade is usually worth making. It is not free, and its cost is no defense of the model.

What the handovers do not carry

None of the handovers is a total victory. Active inference has the mathematics and, in our reading, no demonstrated route to the scale of a polity. The formalism is well posed for an organism or an agent; the step from there to an institution with many semi-autonomous parts has not been made in a way anyone has been able to check. Ostrom has the evidence base, and Chapter 14, *The Evidence: Ostrom*, found its support weakest at exactly the multi-level claim we need. The nested-enterprise principle carries the least weight in the corpus,[2] which is unfortunate, since nesting is the part of her work that speaks most directly to us. Panarchy names transformation and phase, and gives the cycle a vocabulary that most management thinking lacks entirely. Identifying the present phase without hindsight remains difficult, which limits what a practitioner can do with it while the system is still in motion.[3] Bateson, Vickers and Luhmann make meaning unavoidable: after them, no one can treat an organization's account of itself as a neutral report on its environment. In our reading, they do not supply a shared operational test for which meaning should govern. Ulrich makes the boundary political and gives a procedure for questioning it that anyone can run. He does not thereby diagnose the organization inside the boundary, and does not claim to.

Each of them does the job handed over better than we do, and none quietly becomes the whole instrument. The pattern holds across all five. The neighboring literatures are strong where they are specific, and they are not offering, and have not offered, the general claim the tradition made. We do not hold that against them. It marks out the space that remains.

No mathematics, and nothing to run

The most serious of the eight defects catalogued in Part One is also the simplest to state. Variety is not measurable for any real organization, and

[2] Michael Cox, Gwen Arnold, and Sergio Villamayor-Tomás, "A Review of Design Principles for Community-Based Natural Resource Management," *Ecology and Society* 15, no. 4 (2010): art. 38, https://doi.org/10.5751/ES-03704-150438, table 3: nested enterprises has the lowest support ratio of the eleven principles and is the only one whose association with success does not reach significance at the five percent level. Chapter 14 gives the figures and the limits of reading them as evidence about recursion.

[3] Chapter 5, *The Ladder Without a Ruler*, on phases that are easier to identify after the fact than in the present tense; Chapter 11, *The Nesting: Simon, Koestler, Jaques*, on the 2022 systematic review by Rocha and colleagues ("a lack of methods to measure or observe it in real social-ecological systems") and on zu Castell and Schrenk's 2020 method for identifying the phases, which depends on the analyst's choice of parameters.

no one has made it so. Beer said so himself. Having shown how fast variety proliferates, he asked who could use such a measure in a real situation: "Could we ever say: *this* is the variety of the chemical industry, say? Of course we could not." The use of the measure, he concluded, was "not the naive use of counting states, but of matching state generators."[4] That moved the problem without closing it. The model's rules are balances between quantities its author says cannot be counted, and he gave no procedure for checking a balance of that kind. It is a fifty-year absence at the center of a model whose central term is a measure.

Beer's own instruments were thin. The actuality, capability and potentiality triple is a ratio scheme that needs a prior decision about what counts as capability. The statistical filters of Cyberstride were built to detect departures from expectation in a small set of production indices, under severe constraints of time and hardware.[5] Beer described the suite more broadly, as "statistical filtration for all homeostatic loops at all levels of recursion." Neither was ever a variety metric, and neither claimed to be. Most first instruments fall short of the ambition. What matters here is that nothing since has replaced them. The tradition has produced diagrams, case narratives and diagnostic conversations. It has not produced an instrument.

Set beside the neighbors, this bites. System dynamics gives you a model you can execute and be wrong with in public, and that is the property that matters most: a stock-and-flow model makes a commitment the world can contradict. Active inference gives you one you can compute. Aubin's viability theory gives you viability kernels, with the same word in the title and none of Beer's ancestry, an independent formalization of the question of which states permit continued existence under constraint. Powers's perceptual control theory gives you nested control loops that produce numbers, and a hierarchy arrived at from experimental work on behavior rather

[4]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 86–87.

[5]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 163–64, for actuality, capability and potentiality, and 261–64 and 270 for the statistical filters of Cyberstride (Beer's description at 270, checked 29 September 2026 in the calibre PDF); Eden Medina, *Cybernetic Revolutionaries: Technology and Politics in Allende's Chile* (MIT Press, 2011), 70–71, 78–80, 103 and 183–84, for what Cyberstride was, what it ran on and what it produced.

than from a mapping of the nervous system that Beer claimed was homomorphic[6] and never tested.[7] The tradition after Beer kept the diagram and dropped the models Beer attached to System Four.[8]

What the diagram wins in return is real, and it should be claimed openly rather than smuggled in. It can be used in a room, on a wall, by people who will never run a model, and half of what it does is done there. A model that requires a modeler requires one who is present, funded and trusted. A diagram that a management team can argue over on a Thursday afternoon reaches situations no executable model ever will. That advantage deserves to be stated directly. It is no evidence for claims about necessity or measurement. The trouble is not claiming the advantage; it starts when we claim it in one breath and the mathematics in the next.

The residue

What remains after the handovers is short. Five items survive. Part Five is meant to turn each into a claim against which evidence can be registered. Chapter 28, *What Would Count Against It?*, does so for the first and records that the other four are so far exposed only indirectly. Chapter 29, *The Two Cheapest Decisive Experiments*, prices the first tests, on recursion levels and on prediction, and a re-analysis of existing data that bears on the first item.

1. **The closed-set necessity claim.** Chapter 24, *Six Things Nothing Else Does*, argued that five functions are necessary for viability rather than merely common. Beer's form of the claim does not predict that an organization lacking one will fail. He holds that a persisting organization always performs the five in some form, System Two because it "is always present, or the organization would shake itself to pieces," and that a weak function produces pathology, such as

[6]"Models are more than analogies: they are meant to disclose the key structure of the system under study" (Beer, *Brain of the Firm*, 2nd ed., 75); "There can be no interest in analogy at this point" (87). In *Decision and Control* (John Wiley & Sons, 1966), 113, "A scientific model is a homomorphism on to which two different situations are mapped." Checked 29 September 2026 in the calibre PDFs. That the homomorphism was never tested is our reading.

[7]Jay W. Forrester, *Industrial Dynamics* (MIT Press, 1961), and John D. Sterman, *Business Dynamics: Systems Thinking and Modeling for a Complex World* (Irwin/McGraw-Hill, 2000), for system dynamics; Jean-Pierre Aubin, *Viability Theory* (Birkhäuser, 1991); William T. Powers, *Behavior: The Control of Perception* (Aldine, 1973). For active inference, Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo, "Active Inference: A Process Theory," *Neural Computation* 29, no. 1 (2017): 1–49, https://doi.org/10.1162/NECO_a_00912, and Thomas Parr, Giovanni Pezzulo, and Karl J. Friston, *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (MIT Press, 2022), as cited in Chapter 12. The claim is about the existence of an executable model, not about the quality of any particular fit.

[8]Beer's Checo and the quantified System Four model (*Brain of the Firm*, ch. 13) are a lost part of the program, not an absence. The Checo team's experiments were undertaken "in the spirit of a prolegomenon to a full-bodied System Four," and Checo "would be updated every day by the output from Systems 1-2-3, and would promptly re-run a ten-year simulation; and this has never been done" (Beer, *Brain of the Firm*, 2nd ed., 267–68). The System Four model of chapter 13 has five control parameters, A, B, C, X and Y, which a management group would set with knobs in the control center (187–97). Checked 29 September 2026 in the calibre PDF.

the institution without a working System Four that "lives on" as a decerebrate cat.[9] Stated that way, no surviving organization can count against it, which is why Part Five tests it with the coding in Chapter 28 and with viability, not bare survival, as the outcome. No neighboring literature makes the closure claim in this form. The claim remains scarcely tested. We have cases where the functions were present and the organization survived, which is compatible with necessity and with several other explanations.

2. **The transducer.** Chapter 25, *The Transducer*, located regulatory loss at the translation across a boundary, rather than in the channel or at either end. That is a specific and unusual place to put it, and it suggests where to look when the signal that arrived is not the signal that was sent. The principle is Beer's Third Principle of Organization, which Chapter 25 cites from *The Heart of Enterprise*, and it is unmeasured. We have not established a case in which the loss was located there by measurement rather than by attribution after the fact.
3. **The algedonic architecture, with its justification.** Chapter 24 kept an explicit bypass for pain, and for a reason stronger than design preference: hierarchy attenuates precisely the signals that indict it, so a channel that reports upward through the levels being reported on cannot be relied upon. Chapter 4, *Channels Without Owners*, showed that Beer specifies how a signal qualifies, by a statistical filter on points of stability whose criteria the metasystem sets, and leaves the rest open. Who chooses those points, who may raise or contest a signal, and who audits the filter are questions the model poses and does not answer.
4. **One currency across scales.** Variety lets the same diagnosis travel between a person and a planet, and no other framework on the shelf offers a term intended to hold at both ends. Chapter 24 credits the property and concedes in the same section that it may still be metaphor dressed as measure. Until the measurement problem above is addressed, the travel across scales is a property of the vocabulary, and not demonstrably a property of the world.
5. **Freedom as a quantity commensurable with the diagnosis.** Chapter 26, *Freedom as a Quantity*, followed Beer in putting the room left to act after institutional attenuation into the same units as the diagnosis, instead of treating it as a separate normative concern. In *The Heart of Enterprise* he declared it "in principle a computable function of systemic purpose as perceived"; nobody has computed it.[10] That is what makes the tradition's political claim continuous with its technical one. *Residual variety* is our name for the quantity, in the sense defined in Chapter 26 and not in Beer's Third Axiom sense.

[9] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 66; Beer, *Heart of Enterprise*, 143 ("System Four exists — it has to: but it is disseminated") and 265.

[10] Beer, *Heart of Enterprise*, 158.

Part Two identified no other home for this residue. That does not prove no home exists anywhere on the shelf. It is the remainder our comparative search must now put at risk.

Suppose the residue is empty: the closed set only a checklist, transduction only another name for boundary work, the algedonic bypass no more than escalation, variety no common currency, and freedom no quantity. Then the honest conclusion is that the tradition should be wound up. We would rather find that out than preserve it by lowering the test.

Part V

What's at Risk in the VSM

> **Back to School** Judy goes back to school, not to be certified but to be corrected. The lessons are unwelcome and better than she expected: friends in rooms she had written off, colleagues doing her own work more rigorously than she has, coaches whose corrections she had not stayed long enough to hear. She comes out heavier, harder to knock down, more attached to boxing than when she thought she already knew it. What she wants now is a fight she could lose.

Most chapters in Parts One to Three, and Chapter 25 in Part Four, close their charge with a receipt: the smallest piece of work that would settle the matter, what it would cost, and what a negative result would establish. The convention has held since Chapter 1, so its form here will surprise no one. Written one at a time, the receipts could be read as a list of complaints with prices attached. Part Five collects them, orders them by what they would tell us per unit of effort, and makes them into a program.

It is short, dry and priced. It contains only research earned by charges actually made in this book. Nothing is proposed because it would be interesting; the wider program of work this tradition needs stays in the companion register, where it can be argued on its own terms.

Chapter 28, *What Would Count Against It?*, states what observation would make the model give something up. Chapter 9 says that repair is available to us and that we have not made it, and until it is made, every other proposal in this part is decoration. Chapter 29, *The Two Cheapest Decisive Experiments*, sets out the two decisive tests that could be run now, with existing data and existing access, by people already in the field. Chapter 30, *The Bill: A Book Plus a Live Register*, sets out what the whole program costs, and why no total is struck.

The standard is Beer's. "So remember: a model is neither true nor false: it is more or less useful," he wrote in 1985, and in 1966 he had said how usefulness is judged, by "testing of the predictive value" of the model: "If this predictive value is low, we had better start again."[11] The tradition kept the first sentence and dropped the test.

We designed the experiments so that we could lose them, and that is the specification. Each is set up so that a result running against the model is as interpretable as one running for it, and we have said in advance which is which. A negative outcome is then a publishable finding rather than a failure. It is also the only kind of outcome that would change what any of us tells a client next year. Advocacy can be written in the most measured prose available and still leave nothing at stake, so the difference between it and an agenda of scrutiny is not one of tone. What is proposed here leaves something at stake, and the something is ours.

[11] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2, where the sentence is set as three broken lines; Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 113.

Chapter 28

What Would Count Against It?

Chapter 9 closed Part One with the charge that none of the eight defects has ever cost us anything, and it gave the reason: nothing in the model's core, the necessity of the five functions and the recursion claim, has been stated in a form that could be disconfirmed. It also named the repair available to us and noted that we have not made it. Here we make it. The accusations of Parts One to Four become claims that can lose, and for each we state the observation that would oblige us to give something up.

The rule for the conversion binds us before anyone else. A failed test cannot be redescribed after the fact as evidence of some other pathology while the original claim stays unchanged.[1] If the result changes what we say we meant, the registered claim changes with it. A framework that survives by renaming every loss has not survived a test.

The temptation the rule is written against is specific, and it is ours. The model's vocabulary is rich enough to describe any outcome once the outcome is known, so a disappointing result can always be absorbed: the system in focus was misidentified, the recursion level was drawn in the wrong place, the client did not implement what the diagnosis recommended. Any of those may on occasion be true. They are disqualified as a standing defense because they are available every time, at no cost, and after the fact. The conditions below therefore close their escape routes in advance. Where a boundary judgment could later be invoked to explain a loss, it has to be registered before the result is known, or it does not count.

Three kinds of exposure follow, in ascending order of what a loss would cost us. Claims about what Beer wrote can be settled by reading, and losing one costs a chapter. Claims that a VSM diagnosis outperforms its neighbors can be settled by comparison, and losing one costs the practice its warrant for charging as a theory. The claim that the five systems are necessary can be settled by a single well-documented case, and losing it costs the model its central proposition. The cheap tests come first because they are the ones we have least excuse for not having run.

[1]The rule is the one the preregistration literature argues for: Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor, "The Preregistration Revolution," *Proceedings of the National Academy of Sciences* 115, no. 11 (2018): 2600–2606, https://doi.org/10.1073/pnas.1708274114, on the difference between prediction and postdiction, and on hindsight bias as the reason the two are confused unless the questions and the analysis are fixed before the outcomes are seen. Read in the authors' version on OSF (https://osf.io/2dxu5); the published abstract is the same. Chapters 29 and 30 refer to this note.

Only charges traceable to a specific argument in this book belong here. Wider questions, however interesting, stay in the companion register. Normative claims take a different route again: evidence can expose their costs and consequences, but cannot settle what ought to be valued. They return in the Coda.

The three sections share one form. A claim we have made is named; the evidence that bears on it is identified; the result that would count against us is stated, and what we would surrender with it; and the smallest fair test is specified, so that nobody has to fund a decade of fieldwork to hold us to account. The fourth column carries the argument. A condition of disconfirmation that costs nothing when it is met is not a condition at all.

Textual exposure

The cheapest claims to lose are claims about what Beer wrote. Chapters 1 to 3 accuse the tradition of teaching a neurological mapping Beer did not write, of shrinking System Four to environmental scanning, and of citing a theorem more strongly than its proof permits. Each can be tested against a defined corpus, edition and search protocol, and each can be tested by somebody who thinks we are wrong.

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
Current teaching has displaced Beer's anatomical mapping	Reliably paginated *Brain of the Firm* and dated teaching sources	Beer's assignments match the mapping we say is later	Surrender the displacement charge; retain only the modern-neuroscience audit	Two independent coders extract every System Two–Four anatomical assignment from the same edition and record Beer's conflicting assignments as conflicts

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
System Four in Beer does more than scan	Beer's System Four passages, including those Chapter 2 reads as a checkable wager on the System 3–System 4 homeostat[2]	The wider modeling and anticipation functions cannot be found in the stated corpus	Revise the chapter to a charge against later reconstruction, not Beer's text	Register terms and functional criteria before searching the named Beer corpus
The good-regulator result is routinely cited beyond what it proves	The 1970 paper[3] and a dated sample of VSM citations, with Beer's own restatements as a separate stratum[4]	The sample consistently states the paper's actual criterion and concessions	Withdraw the citation-practice charge while retaining the separate question of organizational modeling	Blind-code a reproducible citation sample against declared overclaim criteria

These tests are worth nothing without exact editions, dates and search surfaces. A term missing from an index may still be present in an argument, and a gap in one corpus says nothing about every text Beer wrote. The discipline cuts both ways. A coder who cannot find a functional criterion because we wrote the criterion loosely has learned something about our protocol, and nothing about Beer.

[2]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 182–83 and 189, and Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 240, the passages Chapter 2 relies on. The phrase "prediction channel", which Chapter 2 used for the line to command until 12 October 2026, occurs in none of the eight Beer books searched there; Chapter 2 now reads the same passages as a checkable wager on the System 3–System 4 homeostat.

[3]Roger C. Conant and W. Ross Ashby, "Every Good Regulator of a System Must Be a Model of That System," *International Journal of Systems Science* 1, no. 2 (1970): 89–97, https://doi.org/10.1080/00207727008920220. Chapter 3 sets out what the paper proves and what its authors conceded.

[4]Beer, *Heart of Enterprise*, 235: "any would-be regulator must contain a model of that which needs to be regulated"; and 352: "Every regulator contains a model of whatever is regulated. This model must have requisite variety."

When an analogy bears weight

Chapter 2 asked this chapter for a distinction it needed and did not have: between an analogy that decorates an argument and one that carries load. The neurophysiological charge is the place to draw it, because it has an obvious worthless version, and that is the version most people would write.

"Beer's neuroscience is outdated, therefore the VSM is unfounded" is a bad argument, and it attacks a claim Beer disclaimed in print. In Chapter 10 of *Brain of the Firm* he anticipates exactly the objection that he conflates pain circuits, synaptic biochemistry and sleep-arousal control, concedes it, and states his warrant: "I want to distinguish sharply between the way in which the body achieves results and the logic of the results it achieves. Neurophysiological descriptions in this book are there because they are interesting, and we can learn from them." He gives no glossary of neurophysiological terms because they are "names for bits and pieces inside the body rather than for definable concepts". The training plan drawn up for Chile would teach the cybernetic principles "but not the neurophysiology". Beer also wrote that the more the neurophysiological model is studied in detail, "the more astonishingly accurate does it seem to be in reflecting what happens in the firm." A paper that ignores all this refutes a position Beer disclaimed and only half asserted.[5]

The easy debunking also aims at things Beer never used. *Brain of the Firm* nowhere mentions MacLean, "triune" or "limbic", and it has no Sherrington, no Pavlov, no Moruzzi.[6] The triune brain is discredited in comparative neurobiology and still widespread outside it, in psychology textbooks among other places.[7] It is not Beer's, and attacking it here attacks a straw Beer.

The discriminator therefore has two halves, and it applies beyond this case. An analogy is decorative where the argument survives its removal, and load-bearing where a conclusion is inferred from a property of the source domain. The audit is worth funding only at joints of the second kind, and there are at least three. First, System Four's core mechanism in *Brain of the Firm* rests substantively, not decoratively, on Kilmer, Blum and McCulloch's 1969 computational model of the reticular formation as a behavioral-mode selector,[8] and that model is obsolete. The modern arousal

[5]Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981): the objection and the warrant, ch. 10, "The Biggest Switch", 139–40; the glossary note, 401; the Chilean training plan, ch. 18, "The October Watershed", 317; the claim of accuracy, 109, followed there by the example of the sympathetic ganglia.

[6]Full-text search of an OCR copy of the 2nd edition (Wiley, 1981), including the bibliography, for "MacLean", "triune", "limbic", "Sherrington", "Pavlov" and "Moruzzi", 23 September 2026: no hits. The 1972 first edition has not been searched.

[7]Joseph Cesario, David J. Johnson, and Heather L. Eisthen, "Your Brain Is Not an Onion With a Tiny Reptile Inside," *Current Directions in Psychological Science* 29, no. 3 (2020): 255–60, https://doi.org/10.1177/0963721420917687, who document the layered-brain misconception in introductory psychology textbooks and in psychological research. We have not surveyed management writing for it.

[8]W. L. Kilmer, J. Blum, and W. S. McCulloch, *The Reticular Formation* (Air Force Office of Scientific Research, 1969), the report Beer lists in *Brain of the Firm*, 2nd ed., 407, under

picture runs parabrachial to basal forebrain, with the monoamines modulating arousal rather than generating it, noradrenaline among them as a gain control, and thalamic lesions have remarkably little effect on wakefulness.[9] *The Heart of Enterprise* derives System Four without the reticular formation, from the requirement that it contain a model of the whole system, so a result against the 1969 model would reach the System Four of *Brain of the Firm* and leave that derivation standing.[10] Second, some of Beer's structural inferences draw their force from anatomical position. To argue that a staff function is really command because the diencephalon sits on the vertical axis is to argue from where a thing is. Third, the cortical apex. The cortical placement of System Five is the element current work challenges most directly, and the specialist objection to locating a self anywhere is precisely that it reproduces the homunculus. Beer's own account of identity does not rest on the placement. Viable systems, he wrote, "are self-referential. Their logic closes-in on itself," and in that lies "the maintenance of identity"; System Five he called "the multinode," "an elaborately interactive assemblage of elements."[11] The audit should test whether that distributed account survives, not the apex.

The other half constrains the result rather than the target. We must equally resist writing "the VSM anticipated predictive processing". A triumphalist paper would be worse than a debunking: it would be unfalsifiable, and it would let us count a correction as a confirmation. The test is whether the audit produces at least one statement of the form *this VSM claim is wrong, and here is what replaces it*. An audit that produces none has told us only that we are fluent, which we already knew.

"On 'arousal' mechanisms (System Four)". The journal version gives the authors in a different order: W. L. Kilmer, W. S. McCulloch, and J. Blum, "A Model of the Vertebrate Central Command System," *International Journal of Man-Machine Studies* 1, no. 3 (1969): 279–309, https://doi.org/10.1016/S0020-7373(69)80025-8. Beer calls the ascending reticular formation "the core of System Four" (*Brain of the Firm*, 148).

[9] Patrick M. Fuller, David Sherman, Nigel P. Pedersen, Clifford B. Saper, and Jun Lu, "Reassessment of the Structural Basis of the Ascending Arousal System," *Journal of Comparative Neurology* 519, no. 5 (2011): 933–56, https://doi.org/10.1002/cne.22559: extensive thalamic lesions had little effect on EEG or behavioral measures of wakefulness; large basal forebrain lesions, and lesions of the parabrachial nucleus and precoeruleus area that project to it, left animals unresponsive; and cell-specific lesions of the monoaminergic and cholinergic groups have never reproduced that coma. For noradrenaline as a gain control, Gary Aston-Jones and Jonathan D. Cohen, "An Integrative Theory of Locus Coeruleus-Norepinephrine Function: Adaptive Gain and Optimal Performance," *Annual Review of Neuroscience* 28 (2005): 403–50, https://doi.org/10.1146/annurev.neuro.28.061604.135709.

[10] Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 235: "System Four must contain a model of the viable system of which it is the System Four." The book does not mention the reticular formation (full-text search of the OCR copy, 29 September 2026).

[11] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 4; Beer, *Heart of Enterprise*, 260, where closure is "a self-referential process"; Beer, *Brain of the Firm*, 2nd ed., 201, for the cortex as System Five, and 205, for the multinode.

Comparative exposure

The harder question is whether a VSM diagnosis adds anything beyond neighboring instruments. Part Two credited those literatures with finding the same components independently, and that credit has a consequence. A vocabulary may redescribe an organization elegantly without improving prediction, discrimination or intervention, and elegance is not a result we can bank.

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
A VSM diagnosis identifies recursion levels rather than drawing them	Interaction-rate blocks, time-span measures and blind VSM diagnoses on the same organization	Diagnosed levels fail to track either independent measure beyond chance or analyst discretion	Replace the claim of discovered levels with an explicitly negotiated design convention	One organization, one interaction matrix, one prior diagnosis and two blind analysts
VSM diagnosis predicts a failure or outcome unavailable to simpler alternatives	Pre-outcome engagement documents, outcomes and matched baseline frameworks	VSM coding adds no predictive or diagnostic value over a simple baseline and a neighbor	Retain it, at most, as a facilitation language rather than a predictive theory	Ten retrospective interventions, failures retained and coded by a non-participant
External diagnosis changes the client system in the way the diagnosis recommends	Pre-stated intervention criteria and client as well as consultant outcome accounts	Diagnosis produces no discriminable change, or only changes the consultant's description	Surrender the intervention claim; retain perturbation as the narrower effect	Register outcomes against criteria written before the intervention

None of the three rows asks the model to be right about everything. Each asks only that a diagnosis produced under the model be distinguishable, by somebody who did not produce it, from one produced without it. The bar is low, and we have not yet cleared it. Chapter 29 gives the protocols. What matters here is the consequence of losing. The result cannot be absorbed by declaring that the organization had misidentified its system

in focus, unless that boundary claim was itself registered in advance. Nor can it be absorbed by observing that the client failed to implement, unless implementation criteria were written down before the engagement closed. The second and third rows are where a negative result costs most, because a facilitation language and a predictive theory are priced differently, and much of our practice has been sold at the second price.

Structural exposure

Part Four states the narrow claim and the residue that has no other home, and Chapter 24 sets out six kinds of work that nothing else supplies. The strongest version of that position is a necessity claim, and a necessity claim should face the strongest counterexample available, not the most convenient one.

Beer answered the obvious counterexample before anyone raised it. System Two, he wrote in 1985, "is always present, or the organization would shake itself to pieces." Six years earlier: "System Four exists — it has to: but it is disseminated, and lacks integration."[12] In *Brain of the Firm* he complained that "our age has produced so many large-scale organizations without a System Two," and in the next sentence reported that most of the successful ones he had observed "have been entirely unofficial and largely unrecognized."[13] Where System Four is virtually empty, the metasystem "collapses into System Three," and the institution becomes a decerebrate cat: "It lives on; its viability is ensured by a bogus environment, and is sustained by artificial sustenance."[14] And when someone asks whether an organization that exists is "by definition a viable system," Beer's reply separates the two: "The fact that it is there does not prove that it is effectively there."[15]

Read together, these passages close the obvious route to a counterexample. A persisting organization that lacks a function officially is found to perform it unofficially, and one that performs it badly is diagnosed as pathological, not as a system that has survived without it. Nobody added that defense after a failed test. It is in the founder's statement of the claim, and it is why a surviving system that lacks one of the five is so hard to report. The necessity rows below therefore rest on one coding and one outcome, both registered before any case is read. Other chapters refer to them as the coding in Chapter 28.

The coding. For a named system at a named level and over a stated observation window, each of the five functions receives one of three codes, by indicators registered in advance:

[12] Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 66; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 143.

[13] Beer, *Brain of the Firm*, 2nd ed., 175. See also 183: "there is always a System Four, even if it is not identified in quite the form specified here."

[14] Beer, *Heart of Enterprise*, 265.

[15] Stafford Beer, *Think Before You Think: Social Complexity and Knowledge of Knowing*, ed. David Whittaker (Wavestone Press, 2009), 148.

- *Official*: the function is performed by a role or unit designated to perform it.
- *Disseminated*: the function is performed, but unofficially or by no designated role. This is Beer's disseminated System Four and his unrecognized System Two.
- *Absent*: nobody performs the function in the observation window, as judged by the registered indicators.

The indicators state, for each function, what activity in the records counts as performing it. A coder who finds a function only by reading it into activity after the outcome is known has not coded it. Disseminated is a code, not a refuge: where Beer names the pathology that a weak function produces, as with the "uncontrollable oscillation" he expects without System Two,[16] that expectation is registered with the code.

The outcome. The outcome is Beer's viability, not bare persistence. He took the definition from the dictionary, "able to maintain a separate existence,"[17] and then refused to let the fact of existence stand for it, as the decerebrate cat shows. The outcome indicators are therefore registered separately from the five codes and must not be computed from them, or the test is circular. We take five from Beer's own distinction between being there and being viable, each registered for the named system at the named level before the window opens:[18]

- *O1, separate existence*: at the end of the window the unit is still an identifiable unit with its own decision rights, not dissolved, merged, taken over or run directly by the containing system.
- *O2, effectiveness*: attainment of the purpose the unit itself named before the window, against its own baseline.
- *O3, efficiency*: resources used per unit of that attainment, against the unit's baseline and comparable units, dependence on subsidy included.
- *O4, recovery*: after a change in its environment named in advance, O2 returns to a stated level within a stated period.
- *O5, no rescue*: the unit continued without extraordinary support from outside or from the containing system.

[16] Beer, *Heart of Enterprise*, 177: "System Two is logically necessary to any viable system, since without it System One would be unstable — it would go into an uncontrollable oscillation."

[17] Beer, *Diagnosing the System*, 1, quoting the *Oxford English Dictionary*; the same definition opens chapter 5 of *Heart of Enterprise*, 113, and chapter 8, 199: "The ability of a system to maintain a separate existence, which we call viability, depends on a number of necessary conditions which, in sum, will also be sufficient."

[18] Beer, *Diagnosing the System*, 1: an organization "is viable if it can survive in a particular sort of environment," and its existence is separate, "so that it enjoys some kind of autonomy, it cannot survive in a vacuum"; Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 147, where viability means "capable of independent existence," and 148, where an organization that is there is not thereby "effectively there" or "efficiently there," and some operate "at such an enormous cost that they are becoming less and less viable in front of everyone's eyes."

Loss of viability is registered as O1 failing, or two of O2 to O5 failing, and the five are also reported separately. No indicator may be built from the activities used to code a function: O4 is the restored performance, not whether a planning unit produced a response. Outcome coders do not see the function codes, and function coders do not see the outcomes. Cases without outcome data are excluded, not imputed. O1 is Beer's minimum; whether he would count efficiency as part of viability rather than of performance is our judgment, not his wording.

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
Viability in a changing environment requires the five functions	Cases coded official, disseminated or absent for each function by registered indicators, and for viability by indicators independent of that coding	A system coded absent for one function keeps its viability through the window, or loses it from a structural cause the model cannot see with all five coded present and none pathological	Revise necessity into a contingent configuration claim and identify the substitution or omitted cause	Publish the coding protocol and outcome indicators, then seek deviant cases rather than exemplary ones
The six claimed residues are not supplied elsewhere	Direct conceptual and operational comparison with named neighboring theories	A neighbor specifies the same work with equal or greater diagnostic precision	Remove that uniqueness claim while retaining any useful translation	Adversarial paired reading by specialists from both traditions

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
The five systems are jointly effective even if individually insufficient	A cross-case configuration with outcomes fixed independently of the diagnosis	No configuration containing the five predicts viability better than plausible alternatives	Surrender the configural defense rather than moving the outcome criterion	A preregistered qualitative comparative analysis[19] on an open case register
Regulation fails when a system's relaxation time exceeds the interval between shocks	Relaxation times and shock intervals recorded for named loops before the outcome, and viability indicators independent of them	Systems whose recorded relaxation times exceed the interval between shocks keep their viability as often as those whose do not	Drop the condition as a prediction and keep relaxation time as a descriptive parameter	Register relaxation times and shock intervals for the loops of a few named systems, then follow the viability outcome

The first row is the home of the two candidates Chapter 9 named, restated in the terms of the coding: a system coded absent for one of the five that keeps its viability in a changing environment, and a system that loses its viability from a structural cause the model cannot see, with all five coded present, intact and healthy. We accept both conditions in advance, and both are met by observation rather than by argument. Neither is protected by the objection that the analyst drew the boundary in the wrong place. The coding protocol is published before the cases are sought, and cases are sought because they look deviant, not because they look instructive. Michael D. Pfiffner operationalized viability as a whole for 135 organizations, which shows the claim can be measured; with a reflective specification no result could have shown a function dispensable (Chapter 29).[20]

[19] Charles C. Ragin, *The Comparative Method: Moving beyond Qualitative and Quantitative Strategies* (University of California Press, 1987), as in Chapter 15.

[20] Michael Dominik Pfiffner, "System Viability of Organizations and the Aetiology of Organizational Crisis: A Quantitative Assessment of Stafford Beer's Viable System Model" (PhD diss., Utrecht University, 2017), https://hdl.handle.net/1874/356772; the reflective measurement model, 180. Chapter 18 discusses the design.

Beer wanted a test of this kind for his own work. Of the identity exercise in *The Heart of Enterprise*, which clusters managers' answers about the cohesion of their enterprise, he wrote that "maybe (one day) thresholds will be established, to exceed which will predict collapse."[21] Nobody has established them. A registered threshold of that kind would be a structural claim that could lose, and it is the founder's proposal, not ours. The same book supplies a second candidate, the last row of the table above. Once a loop is disturbed, Beer asks how long it takes to return to stability, and answers that "if the system has not returned to stability before another shock drives it once more towards instability, then regulation is impossible."[22] Relaxation time and the interval between shocks can be recorded before the outcome is known, which makes the condition a prediction that could fail.

Chapter 27 lists five items of residue that Part Five must turn into claims against which evidence can be registered: the closed-set necessity claim, the transducer, the algedonic architecture together with the justification offered for it, one currency across scales, and freedom as a quantity, which Beer declared computable in principle and nobody has computed, and which we call *residual variety* in the sense defined in Chapter 26. Only the first is converted above in its own right. The transducer has its row in Chapter 25's receipt, a blind coding of documented regulatory failures that loses if failures do not concentrate at crossings, or if coders cannot tell crossing from channel without VSM labels as their evidence. The algedonic architecture has its row in Chapter 30, where the channel is red-teamed against modern incident and escalation practice. Freedom as a quantity takes the row below. One currency across scales cannot be converted yet: until the measurement defect Chapter 27 names is addressed, there is no quantity whose travel between a person and a planet could fail.

[21] Beer, *Heart of Enterprise*, 432.

[22] Beer, *Heart of Enterprise*, 390. Checked 29 September 2026 in the calibre PDF.

Claim at risk	Evidence that bears	Result against us	What changes	Smallest fair test
Residual variety, in Chapter 26's sense, can be counted	The distinctions available to a named unit, recorded before and after one deployment by a procedure registered in advance	Two coders cannot agree on the count, or the count does not move where the design says room was returned or taken	Withdraw freedom as a quantity and keep it as the question Chapter 26 asks at diagnosis	Count the distinctions available to one unit before and after one deployment, as Chapter 26 proposes

The companion register makes configural analysis, structural disconfirmation and the intervention question answerable across cases. Their claims have been converted above; Chapter 29 keeps the protocols in full, and Chapter 30 prices the whole program. The division is deliberate. It stops the same paragraph from being counted once as a test and again as an experiment, which is the arithmetic by which a research agenda comes to look larger than it is.

Chapter 29

The Two Cheapest Decisive Experiments

Two experiments the accounting produced

Chapter 5, *The Ladder Without a Ruler*, found that the model has no procedure for deciding where a recursion level lies. Chapter 9, *Why None of This Ever Cost Us Anything*, and Chapter 14, *The Evidence: Ostrom*, found that it has no evidence base. Each finding produced an experiment, and these two come first in the program because they are the cheapest decisive work available to us and the least excusably undone, not because they matter most. They came out of the book's own accounting rather than a wish list. Chapter 28 states the conditions under which a result from either would count against the model, and we do not restate them here. Chapter 30 sets out what the whole program costs.

We describe these two in more detail than the rest, because each has to be legible enough that a reader who thinks the design is wrong can say where. For each there is a question, a design, a decision rule fixed before the data is seen, a price where the work we have done supports one, and the negative result the experiment would establish. Where a design choice is contestable we say so and name the alternative. A proposal that hides its choices cannot be criticized, and so cannot be improved.

Experiment one: measure the recursion level

The criterion Chapter 5 showed the model has never had.

The claim under test. Beer's test for whether a candidate is a recursion level is, in our paraphrase, that the contained system must be viable in its own right and one of the units that do what the containing system does.[1] It is a good test, and it measures nothing. Beer claimed more for it: the

[1]The wording is our paraphrase of Beer and not a quotation. His nearest statements: the operational elements of a viable system exhaust its "basic activities, namely those which exist to do what the system does" (Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 116); a division that "could in principle be hived-off as a separate business" is a viable system, and the accounting system that prepares the group accounts is not (ibid., 118); and "the viable system is necessarily a producer of the organization, and not just an adjunct to it, however important" (Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 8).

principle of recursion "is by no means loosely expressed: it is rigorously determined by the Law of Cohesion."[2] What he supplied to apply it are rules of thumb, such as the warning that more than seven or eight operational units means "you may be missing a whole level of recursion."[3] Two competent analysts who disagree about a boundary have no procedure for settling the disagreement beyond those heuristics. The same hole appears one level down, in variety itself. Variety is a count of distinguishable states, and Beer says plainly who does the distinguishing: the observer, relative to the purpose attributed to the system. "Everything will depend on how the system is defined," and two people "will not necessarily agree" on its boundaries.[4] His only arbiter between observers is predictive usefulness, and the model has no procedure for agreement. The experiment asks whether the levels a trained diagnostician draws correspond to any structure recoverable from behavior, or whether the diagnostician produced them.

The design applies three instruments to the same organization, and each is informative whichever way it comes out.

The three instruments, and what could go wrong with them

First, the rate-separation study. Simon defined a level as a separation in interaction rates: parts within a block interact much faster than blocks interact with each other.[5] That is computable from data that did not exist when Beer worked and now exists everywhere, in message logs, ticketing systems, transaction graphs, incident propagation and deployment dependencies. Take a set of organizations that have been diagnosed with the VSM, build the interaction matrix, look for block structure, and ask whether the blocks coincide with the recursion levels the diagnostician drew. Simon's blocks and Beer's recursions are related but not identical objects. Beer's criterion, set out in Chapter 10, is viability, not interaction density, and his System Two exists because the operational units inside one recursion affect each other enough to oscillate without it.[6] Coincidence would therefore be a convergence of two criteria, not a confirmation of one.

The contestable choice is which trace stands as the proxy. Message logs record who talks to whom and not what the talk decides. Ticketing

[2]Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 229; see also Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 307–8, and Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2, on "an absolutely precise definition of viability."

[3]Beer, *Diagnosing the System*, 19; see also 4, on considering "a trio of viable systems at any one time."

[4]Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 247 and 243; Beer, *Diagnosing the System*, 98. On predictive usefulness, *Decision and Control*, 113 and 243.

[5]Herbert A. Simon, "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82. Chapter 11 quotes the near-decomposability argument with pages.

[6]Beer, *Heart of Enterprise*, 177 ("without it System One would be unstable — it would go into an uncontrollable oscillation") and 265 ("the whole metasystem collapses into System Three").

systems record work formalized enough to be ticketed and miss the rest. Deployment dependencies describe a technical architecture that may or may not be the organization's. Each can be argued to measure the wrong thing. The alternative design builds the matrix from two or three independent traces in the same organization and treats agreement between them as a precondition for treating any of them as evidence about the diagnosis. It costs more, and it removes the easiest objection: that a negative result merely reflects a badly chosen data source.

Second, the time-span study, on the same organizations and preferably in different hands. Jaques's time-span of discretion attaches a number to a role, and a stratum follows from it.[7] Measure it at the boundaries the diagnostician drew. If Beer and Jaques are describing the same structure, the boundaries should coincide. Beer warned against exactly that reading: "If the term *recursion* is allowed to become a meaningless appendage to the term *organizational level*, then we do not have a valid model of a viable system."[8] A time-span stratum is a managerial stratum inside one hierarchy; a recursion is a viable system with its own metasystem. On Beer's terms the two need not coincide, and where they do, the diagnostician may have drawn levels of management and called them recursions. The registration has to say in advance how coincidence and non-coincidence will each be read, and we register the rule here. Where time-span strata coincide with diagnosed recursions and the units fail Beer's viability test, the reading is management levels drawn as recursions. Where diagnosed recursions cross strata and the units pass the test, the reading is support. Partial overlap is reported, not scored. The limit is availability. Time-span can be established only where role records or interviews support it, which is why the receipt in Chapter 5 makes this instrument conditional on what the records permit.

Third, the variety instrument. Galbraith's descendants spent four decades building construct-validated measures of information-processing requirements, of information-processing capacity and of the fit between them, and testing them on large samples.[9] Nothing prevents the same being done for the variety claims, except that nobody has done it. A fit measure of that kind suits Beer's own use of variety better than any absolute

[7] Elliott Jaques, *Measurement of Responsibility: A Study of Work, Payment, and Individual Capacity* (Tavistock Publications, 1956), for the measure, and *Requisite Organization: The CEO's Guide to Creative Structure and Leadership* (Cason Hall, 1989), for the strata; Chapter 11 sets out the empirical record.

[8] Beer, *Heart of Enterprise*, 308, where the passage continues: "The levels of recursion in any such organizational model answer to mathematical criteria." Chapter 10 states the criterion in full.

[9] For example Michael L. Tushman and David A. Nadler, "Information Processing as an Integrating Concept in Organizational Design," *Academy of Management Review* 3, no. 3 (1978): 613–24, https://doi.org/10.5465/amr.1978.4305791, and G. Premkumar, K. Ramamurthy, and Carol Stoak Saunders, "Information Processing View of Organizations: An Exploratory Examination of Fit in the Context of Interorganizational Relationships," *Journal of Management Information Systems* 22, no. 1 (2005): 257–94, https://doi.org/10.1080/07421222.2003.11045841, as cited in Chapter 12, which calls them points on a line, not a survey of it.

count. He ruled the count out himself ("Of course we could not") and said the measure's use lay in "matching state generators," not in "the naive use of counting states."[10] This instrument is the least developed of the three and we do not claim it is ready. It would have to be adapted rather than applied.

The decision rule, registered in advance. The block-detection rule is fixed and written down before the interaction data is inspected: what counts as a block, at what threshold, and how a block is matched to a diagnosed level. The rule we propose, as in Chapter 5, is a nested degree-corrected stochastic block model with the number of levels chosen by minimum description length, which leaves no free threshold for the blocks themselves.[11] Two analysts work blind, one holding the diagnosis and the other the matrix, and neither of them makes the comparison afterward. Registration matters more here than the threshold chosen. A rule chosen after the matrix has been seen can be tuned until the blocks agree with the levels, and nobody, the analyst included, will be able to tell afterward whether it was.

Receipt — Experiment one, first stage. One organization, one interaction matrix, one existing diagnosis, two analysts each blind to the other's material, a block-detection rule registered before the data is seen, comparing detected blocks, diagnosed recursion levels and, where records permit, Jaques's time-spans of discretion. Priced in Chapter 5 at two analyst-weeks plus data preparation. Negative result: the diagnosed levels have no proxy in observed interaction structure. That counts against the equivalence of Beer's criterion with Simon's or Jaques's, and against any claim that recursion levels can be recovered from interaction data; it does not by itself count against viability as Beer's criterion. The multi-organization extension — several diagnosed organizations, several traces each, the variety instrument adapted and applied — is the natural continuation and is not yet priced; it should not be commissioned until the single-organization stage has produced a result either way.

What experiment one would settle, and whom it would cost

The design bears on more than the headline question. Where diagnosis and proxy diverge, which one is wrong? The blind design keeps that open rather than settling it, since neither instrument has prior standing over the other. Is Beer's criterion equivalent to Simon's, to Jaques's, to neither, or to

[10] Beer, *Heart of Enterprise*, 86–87: "the potential use of the variety measure is not the naive use of counting states, but of matching state generators."

[11] Brian Karrer and M. E. J. Newman, "Stochastic Blockmodels and Community Structure in Networks," *Physical Review E* 83, no. 1 (2011): 016107; Tiago P. Peixoto, "Hierarchical Block Structures and High-Resolution Model Selection in Large Networks," *Physical Review X* 4, no. 1 (2014): 011047. We have not read either paper in full; the method is to be confirmed with a network methodologist before registration.

both at different scales? And underneath all of them: is there any quantity in this model that two observers can measure independently and agree on?

The worthless version of this study is a paper that proposes a metric and applies it to one sympathetic case. Two things separate the useful version from it. The first is inter-observer agreement: the study is only worth doing if two analysts work blind and their results are compared. The second is the publication of a negative result, on the terms set out under the conditions shared by both experiments. A finding that VSM recursion levels do not track rate separation would be the most informative single result this community could produce.

If the levels do track something, the practical yield is that a diagnostician could show a client the boundaries were not chosen arbitrarily. Disputes about the system in focus would acquire evidence without thereby becoming settleable, since who may draw a boundary is a separate question from what the data supports. And a failed diagnosis could be told apart from a badly bounded one, which at present it cannot.

The audiences already exist. Organizational network analysis has the data and the methods and no theory of levels. The information-processing tradition in information systems and operations management has the instruments. The requisite-organization practitioner network has fifty years of time-span data and a strong interest in independent corroboration. Complex-systems and modularity researchers treat near-decomposability in real organizational data as an open empirical question. And every consultancy selling organizational design would rather have a number than an argument.

As far as our searches found (Crossref and web, September 2026), nobody is working on the joint question.[12] Organizational network analysis measures interaction structure and does not ask about viability. The requisite-organization network measures time-span and does not talk to cyberneticians. The information-processing literature measures fit and has rarely engaged System Four. The three instruments live in separate literatures, and our search has not yet found a published study that puts any two of them on the same organization.

[12]Crossref bibliographic search, 29 September 2026, first eight results for each of "organizational network analysis viable system model", "information processing view viable system model", "time span of discretion viable system model", "interaction matrix recursion viable system model" and "near decomposability organization viable system model"; and one web search the same day combining "viable system model" with the three instruments. The nearest hit, judged from its title, is Gary Preece, Duncan Shaw, and Haruo Hayashi, "Using the Viable System Model (VSM) to Structure Information Processing Complexity in Disaster Response," *European Journal of Operational Research* 224, no. 1 (2013): 209–18, https://doi.org/10.1016/j.ejor.2012.06.032, which is not a study of two instruments on one organization. Manuel E. Sosa, Steven D. Eppinger, and Craig M. Rowles, "The Misalignment of Product Architecture and Organizational Structure in Complex Product Development," *Management Science* 50, no. 12 (2004): 1674–89, https://doi.org/10.1287/mnsc.1040.0289, set product architecture beside organizational structure outside the VSM, as the title reports; we have not read it.

The cost of a result falls on us first. A measurable criterion means a diagnosis can be shown to have been drawn wrong, and our diagnostic authority currently rests on the fact that it cannot. The pedagogy is exposed too, since a criterion that disagreed with the model's own would have to be taught alongside it. And so is the model, if the answer is that nothing recoverable from data corresponds to a recursion level at all.

Experiment two: test whether a VSM diagnosis predicts

The evidence base Chapters 9 and 14 showed the neighbors have and we do not.

The claim under test. Ostrom's field built a coding instrument of four hundred and fifty-five questions, coded eighty-six cases against it, and read fourteen in depth, five of them failures. The result was then independently reviewed and reanalyzed three times, and the reviews reported which of her principles the evidence did not support.[13] The most recent systematic review of the Viable System Model started from seven hundred and sixty-eight Scopus records and forty-five more found by hand, read a hundred and twenty in full, kept twenty-one studies, and concluded that most of what exists is model and simulation rather than implementation.[14] The gap is not one of counting. It is that one field built an instrument, ran it, and printed the findings that went against it. The experiment asks whether a VSM diagnosis predicts anything a simpler account does not.

The register, and the test built on it

The instrument is a retrospective register of VSM interventions, designed to test predictions rather than collect success stories. The register itself is unglamorous and overdue. Each entry records the system in focus and who named it, the boundary judgments made and by whom, the diagnosis, the intervention, the success criteria stated in advance, the outcome, and the client's own account of the outcome alongside the consultant's.

[13] As in Chapter 14: Elinor Ostrom, Arun Agrawal, William Blomquist, Edella Schlager, and S. Y. Tang, *Common-Pool Resources (CPR) Coding Manual* (Workshop in Political Theory and Policy Analysis, Indiana University, 1989), with the figure of 455 questions from Michael J. Bernstein et al., "Mapping Ostrom's Common-Pool Resource Systems Coding Handbook to the Coupled Infrastructure Systems Framework to Enable Comparative Research," *International Journal of the Commons* 13, no. 1 (2019): 528–52, https://doi.org/10.18352/ijc.904; Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge University Press, 1990); and the three reviews, Cox, Arnold and Villamayor-Tomás (2010), Baggio et al. (2016) and Gari et al. (2017), cited in full in Chapter 14. Chapter 14 also records that the manual itself lists eighty-five coded cases, not eighty-six.

[14] Mohammad Reza Fathi, Soraya Birami, Alireza Payvar, and SeyedAli Doorafshan, "A Systematic Review of the Viable System Model: Applications, Insights, and Future Directions," *Journal of Systems Thinking in Practice* 4, no. 3 (2025): 109–45, https://doi.org/10.22067/jstinp.2025.91552.1136. The selection figures are from the PRISMA diagram, 113 (fig. 2); on reliance on simulation rather than implementation, 131 (table) and 139.

The two accounts stay separate on the record rather than being reconciled into one, because the divergence between them is itself a finding.

None of this has to be invented. Ostrom's coding instrument is published and could be adapted. So could the sociotechnical tradition's, which reviewed a hundred and thirty-four experiments in 1982 and said candidly, in the same paper, that most of them had not done the thing the theory claimed.[15] Adapting a published instrument is cheaper, and easier to criticize: a reader can set our version beside the original and see what we dropped.

The decision rule, registered in advance. Diagnoses and predictions are coded from documents created before outcomes were known, by someone who did not conduct the intervention. The VSM claims are then compared with a simple baseline and with at least one neighboring framework, on criteria fixed before the comparison is run. Each case faces the same test, in the same terms for all three: did the diagnosis identify, in advance, a failure mode or outcome that the baseline and the neighboring framework did not? Around that test, the register makes configural analysis, structural disconfirmation and the intervention question possible across cases. Those are the reasons to build a register rather than run a single comparison.

Three choices are contestable. The first is retrospection. Coding from historical engagement documents is fast and cheap, and it depends entirely on what was written down at the time. A prospective design, registering diagnoses and criteria before interventions begin, is stronger and returns its first result years later. We propose the retrospective version because it can start now, and we accept that its findings are weaker in exactly the way that dependence implies. The second is the comparator. A badly chosen baseline makes the model look good, so the baseline has to be named in the registration, not selected once the cases are in hand. The third is breadth against depth. Ostrom's field coded eighty-six cases and read fourteen in depth; a small register can imitate the depth or the breadth but not both, and for the opening stage we have chosen depth.

Receipt — Experiment two, first stage. Ten interventions, retrospectively coded by someone who did not conduct them, against criteria reconstructed from the original engagement documents rather than from memory, and compared with a simple baseline and at least one neighboring framework. Publish the ten, including those that failed and those where nobody wrote the criteria down. That establishes both the instrument and the honesty of the exercise, and it can be done in a year. Negative result: VSM diagnoses do not identify, in advance, failure modes or outcomes that the baseline and the neighboring framework also identify. The register's continuation beyond the first ten entries is not yet priced.

[15] William Pasmore, Carole Francis, Jeffrey Haldeman, and Abraham Shani, "Sociotechnical Systems: A North American Reflection on Empirical Studies of the Seventies," *Human Relations* 35, no. 12 (1982): 1179–1204, https://doi.org/10.1177/001872678203501207, as cited in Chapter 13.

What experiment two would settle, and whom it would cost

The questions in reach are the ones a client asks first. Does a VSM intervention improve anything, against criteria fixed before it starts? Which of the five functions, coded absent or disseminated under the coding in Chapter 28, actually predicts the pathologies Beer says it predicts, such as oscillation where System Two is missing and the collapse of the metasystem into System Three where System Four is empty?[16] Do those pathologies occur in organizations where the model was never applied, which is to ask whether they are findings or artifacts of our vocabulary? What is the failure rate, and what distinguishes the failures from the rest? And is our record unblemished because the model is good, or because it has never been tried in a way that could mark it?

The worthless version here is a case-study collection, which we already have, written by people with an interest in the outcome. A register differs from a collection in two respects that have nothing to do with size: the criteria were fixed before the outcome was known, and the coder had nothing to gain from the answer.

If the answer is favorable, the register buys a response to the first question any serious client, funder or regulator asks, which at present we cannot answer. It gives training a basis for separating what is known from what is believed. And it answers the reasonable charge that a framework absorbing every outcome is a language rather than a theory.

Again the interested parties are mostly outside the tradition. Metaphorum and the practitioner networks, whose credibility depends on this more than is generally admitted. Evidence-based management, which would take the data seriously and the vocabulary not at all. The commons and sociotechnical research communities, which have run this exercise on their own traditions and know how it goes. Journals that publish registered reports. And the model's critics, for whom a register would be the first thing to argue with rather than about.

No register exists. Quantitative tests do: two earlier St. Gallen dissertations, which we have not read (Chapter 18), and the two we examine there. One is Schwaninger and Scheef's 2016 survey.[17] The other, and the nearer, is Michael D. Pfiffner's 2017 dissertation,[18] an instrument rather than a survey: an operationalization of the model's own conditions, applied to 135

[16] Beer, *Heart of Enterprise*, 177 ("without it System One would be unstable — it would go into an uncontrollable oscillation") and 265 ("the whole metasystem collapses into System Three").

[17] Markus Schwaninger and Christine Scheef, "A Test of the Viable System Model: Theoretical Claim vs. Empirical Evidence," *Cybernetics and Systems* 47, no. 7 (2016): 544–69, https://doi.org/10.1080/01969722.2016.1209375. Its hypothesis on System Three was rejected; Chapter 18 discusses the result and how the paper explains it.

[18] Michael Dominik Pfiffner, "System Viability of Organizations and the Aetiology of Organizational Crisis: A Quantitative Assessment of Stafford Beer's Viable System Model" (PhD diss., Utrecht University, 2017), https://hdl.handle.net/1874/356772. Pages below follow the Utrecht repository version, printed pagination.

organizations by third-party raters independent of the organizations' management (22 respondents, five of them consultants or former consultants of the author's firm; one rated 23 of the 135 cases).[19] Neither is a register, since neither records interventions, criteria fixed in advance, or outcomes over time. But the second changes what this chapter should recommend first, and the change is set out below. The 2025 systematic review is an audit of the literature. Its findings (limited empirical validation, few real implementations, almost no longitudinal work) make the case for this direction and cannot substitute for it. The neighboring fields have done the work and published their instruments, which removes the last excuse.

The cost falls on everyone in the tradition who bills for a diagnosis, ourselves included. A register with pre-stated criteria and published failures is an invitation to be measured, and for fifty years we have been able to decline it. In the short run it also puts the model's reputation at risk, because the first honest register will look worse than the case literature it replaces. Every honest register looks like that at the beginning. A community that cannot survive its own first negative result has no claim to be a science.

A cheaper check, if the data can be had

We wrote the pages above before we had read Pfiffner's dissertation. Reading it produced a candidate that costs less than both experiments and answers a question neither of them reaches.

The data already exist. One hundred and thirty-five organizations, seventy-four of which had undergone an existential crisis and sixty-one of which had not; thirty-five independent variables; indicators organized by the model's own functions, with separate blocks for Systems One to Five and for the principles of recursion, requisite variety and communication. The expensive part of an empirical study, finding the cases, reaching the informants and collecting the ratings, was done in Switzerland between December 2014 and April 2015.[20]

The study did not test whether any function is necessary, and Chapter 18 explains why it could not. The measurement model is reflective and specified downward, with viability treated as a single latent factor of which the functions are interchangeable expressions. In that specification the functions are expected to move together. They duly do, and no possible pattern of results could have shown one of them to be dispensable.

[19] Pfiffner, "System Viability of Organizations": "knowledgeable and unbiased third party assessors," unbiased "as compared to managers or owners of the respective organizations," 133 and n. 23; five respondents addressed "among (ex-) consultants of the author's employing firm," 179; twenty-two respondents from sixteen organizations, an average of 6.1 cases each, one of them contributing twenty-three, 181.

[20] Pfiffner, "System Viability of Organizations": data collection, 179; thirty-five independent variables in the sample of 135, and the reflective measurement model, 180; seventy-four crisis and sixty-one non-crisis cases, 183. The final path model drops five of the twenty-two VSM items, among them P.1, upward recursion (200–201), so the block list in a registered rule has to be fixed against the final model and not the questionnaire.

The claim under test. The model asserts that the five systems are necessary: without each, no viability. A reflective model cannot assess that. A formative one can come closer, because it treats the functions as constituting viability rather than expressing it, and then asks what each block contributes that the others do not. Two limits follow from Beer's own statements. He held that "in a viable system all five subsystems are dependent on each other,"[21] so a block that adds little once the others are in the model may be moving with them rather than doing nothing. And Pfiffner's items rate how well each function works, not whether anyone performs it. The re-analysis tests whether each function's measured quality adds discrimination of crisis that the others do not supply. It tests necessity in Beer's sense only for cases in which absence is coded, as in Chapter 28, and the crisis groups are not the same outcome as Beer's viability.

The design. Re-specify the measurement model as formative and re-estimate on the existing data. For each of the five systems and each of the three principles in turn, remove that block and record the change in the model's ability to separate the crisis group from the non-crisis group. Because every run uses the same data, the comparison is internal and clean.

The decision rule, registered in advance. Before re-estimation, and in writing: the drop in discrimination a block must produce for its function to count as carrying weight the others cannot supply, and the threshold below which a function is reported as not distinguishable from the rest of the instrument. Fix which statistic is used and what counts as a difference, and how variance shared between blocks will be read, given Beer's statement that the five are mutually dependent. Publish the rule before the first run.

The contestable choices are real. Re-specifying someone else's model is an interpretive act, and the original author should be a party to it, not its object. The instrument's items were written for a reflective design and some may not carry formative weight; that would itself be a finding, about the instrument rather than the model. The outcome variable inherits the problem Chapter 18 describes, since it contains inverted viability items. The re-analysis should therefore be run both on that construct and on the original crisis and non-crisis group membership, which the dissertation reports but did not use in the final model. The ratings are clustered, about six cases per rater, and the analysis has to allow for that. And the sample is Swiss, from the mid-2010s, and selected at the extremes.

Receipt — the formative re-analysis. Obtain the dataset, re-specify the measurement model as formative, and run the block-removal comparison under a rule published beforehand, on both the composite outcome and the original group membership. Price: one request to the author, and roughly three to four weeks of a statistician competent in structural equation modelling; no fieldwork, no access negotiation, no new instrument. This is cheaper than either experiment above by a wide margin, and it is

[21] Beer, *Diagnosing the System*, 128.

the only route we have found to the necessity claim that does not require collecting a new sample. The re-analysis cannot code a function as absent in the sense of Chapter 28, because the instrument has no such answer: raters scored how far a described ideal matched the organization, on an eight-point scale with "don't know" as the only other option.[22] If a subsystem's removal makes no difference under the registered rule, the claim that the subsystem carries weight the others cannot supply is in trouble — and Schwaninger and Scheef's System Three result, reached from an entirely different instrument and sample, says where to look first. If every block carries weight, the model has its first corroboration of necessity from a design that could have refused it.

What would make this worthless. Running it without the rule fixed in advance. Running it without the author. Reporting only the blocks that behaved. And treating a Swiss sample of 135 as more than it is: one dataset, at the extremes, from one decade.

Conditions shared by both experiments

The two designs share a small set of conditions, and neither is worth commissioning without them. Criteria and exclusions are registered before results are inspected. Analysts work blind wherever the material permits. Failures stay in the dataset. Missing criteria are recorded as findings, not repaired from memory. Data and exclusions are open to anyone who wants to check the analysis. Negative results are submitted for publication.

The last condition is the one that costs something. A finding that diagnosed levels do not track independent measures, or that VSM predictions do not beat the alternatives, must be publishable and must be published; the conditions under which such a finding would count against the model are the ones Chapter 28 has already set out. Neither experiment requires a theoretical advance. Both require the willingness to be seen losing.

[22] Pfiffner, "System Viability of Organizations," 146 and 159.

Chapter 30

The Bill: A Book Plus a Live Register

A bounded audit here, an analyst-week there, a review paper, a facilitated session, an archive fee. Each receipt in this book, read on its own, is small enough to seem costless. Read together they are a program, and a program has to be paid for by people with names, calendars and budgets. Until the items stand beside one another, we can agree with every one of them and commission none. Adding them up is meant to make that harder.

The receipts come from the whole book. Part One set out the defects we inherited and never tested, and closed on the charge that nothing in the model's core, the necessity of the five functions and the recursion claim, has been stated in a form that could be disconfirmed. Part Two credited the neighboring literatures that arrived at the same components without us. Part Four narrowed the claim to what we can defend and named the residue that has no other home. Part Five has been turning those charges into work: experiments with a price on them, each with a statement of what a negative result would establish. What remains is the total.

The bill has three parts, and only the first is money. The first is the work: what would have to be done, and what it would cost at minimum. The second is the people: which competences the work requires, most of which are not ours, and some of whose holders have every reason to decline. The third is the standing record: the obligation to keep the account open after the book is printed, so that items close on evidence rather than on fatigue.

Why the complete list is not printed here

The complete research program is not here. These pages print the bill earned by *Punch and Judy*; the wider program belongs in a maintained, dated and versioned companion register, the live register of this chapter's title.

The accounting made that decision, not convenience. One of this book's sharpest findings, in Chapters 19 and 29, is that the field never built and maintained a record of what it did. Beer asked for one: a small team of "organizational pathologists" to study organizations that are "viable no more," and thresholds for his identity exercise "to exceed which will

predict collapse."[1] A numbered list, fixed on the day of typesetting and never touched again, would repeat in the answer the failure named in the question. A register is not merely a publishing format. It is the thing the accounting demands, and building it is the first item on it.

There are three further reasons. The accounting concerns more than fifty years that have already happened, and it does not decay. An agenda is a claim about what is not yet known, and it decays the moment somebody does the work, so the two do not belong in one object. If this book succeeds, problems close; a printed chapter cannot record its own success, while a register can show a problem closed, by whom, and with what result. And a register is the only part of this book a reader can act on rather than assent to. A problem can be claimed, a case added, a verdict contested, and none of that is available on a page.

One consequence, stated once: this book never refers to a research problem by number. Numbers belong to the register, and they will change as items are added, closed, split and merged. The printed bill names the claim and the work. The register is kept by Krishan Mathis and revised quarterly, and each version is deposited on Zenodo with its own DOI, so that a citation of the register names a fixed state of it.

What would make the register worthless

Splitting the bill from the register helps only if the register meets a standard the tradition has not held itself to before. As already said in Chapter 29, the worthless version is a case-study collection, which we already have, written by people with an interest in the outcome. The discriminators are three, and strict: success criteria registered before the intervention;[2] failures included; and the register open, so that a hostile reader can recount.

None of the three is expensive. Each is uncomfortable, and that is what makes them work. Registration removes the option of discovering the criterion after the result. Including failures removes the selection that made the existing literature persuasive. Openness removes our privilege of being the only party able to check the arithmetic. Anything short of that reproduces the existing literature at greater length. These conditions apply to every item below. The conditions specific to individual experiments were set out in Chapter 29 and are not repeated here.

The bill, ordered by minimum viable cost

The costs are bands rather than quotations: **desk** means weeks of documentary work; **small study** means one organization or a bounded corpus; **program** means multiple cases, specialist collaboration or longitudinal

[1] Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 150; Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 432.

[2] On registration before the result, see the note on Nosek et al. (2018) in Chapter 28.

access. Ordering by cost makes refusal visible. The cheapest items do not wait for the expensive ones.

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The mapping taught is not the mapping Beer wrote	Audit Beer's System Two–Four assignments against dated teaching sources, recording Beer's own conflicting assignments as conflicts[3]	Paginated Beer text; captured primers; independent extraction	Desk	The assignments materially match	The tradition drifted from its source	History-of-cybernetics researcher plus neuroanatomist	Neurophysiological remapping

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The good-regulator theorem is cited beyond its proof	Blind-code a reproducible citation sample	Search protocol, corpus, coding rules, second coder; Beer's own restatements as a separate stratum[4]	Desk	Citations consistently preserve the criterion and concessions	Citation practice substitutes authority for reading	Information theorist and research assistant	Foundations and citation audit
System Four has been reduced to scanning	Search current teaching material against registered functional criteria; the Beer side is answered[5]	Editions, search surface, quotations and dated captures	Desk	Teaching retains the wider function	A central function disappeared in transmission	VSM scholar and independent textual coder	System Four dependency analysis

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The six components have independent lineages, and the related entries do not stand in for recursion	Completeness and priority audit: earliest independent statement of each component, citations checked in both directions, and for each entry whether it reached Beer's object or a related one (Chapter 10)	Six bounded searches; preregistered evidence table with a related-object column; external spot-checks	Desk	One or more components has no independent lineage	The model's components are non-arbitrary	History-of-science researcher and an external spot-checker	New — added 29 September 2026

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The method literature dropped Beer's freedom quantity	Systematic survey of the VSM method literature (Chapter 26): stated inclusion rule; coding of freedom and management vocabulary and of where authority over the model is placed; pilot from Espejo and Reyes (2011), then Hoverstadt and Martin Pfiffner (2020	Inclusion rule, corpus, coding scheme, reported counts, second coder	Desk	The method literature keeps freedom as a quantity, or asks whose residual variety an intervention should increase	The method descends from the handbook's narrowing	VSM scholar plus a second coder	New — added 29 September 2026

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
An algedonic channel performs a distinct safety function	Red-team the mechanism against modern incident and escalation practice	Failure cases, escalation traces, comparator mechanisms	Desk / small study	No discriminable function or advantage remains	The channel is structurally necessary rather than metaphorical	Safety researcher and incident-response team	Algedonic red team
Recursion levels are discovered rather than drawn	Compare blind VSM diagnosis with rate separation[6] and timespan measures	Interaction matrix, prior diagnosis, registered protocol	Small study	Levels fail to track either independent measure	The recursion ladder has an observer-independent rung	Organizational network analyst, VSM analyst, requisite-organization practitioner	Recursion measurement protocol

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
VSM diagnosis predicts something alternatives miss	Reconstruct pre-outcome predictions for ten interventions and compare them with a baseline and a neighbor	Original engagement documents, outcomes, failures, independent coding	Small study	VSM adds no predictive or diagnostic value	The framework is more than a redescription language	Evidence-based-management team and independent coders	Intervention register pilot
The embodied observer can be specified without smuggling in a view from nowhere	Write the 4E reconstruction and test its operational distinctions on one diagnosis	Sensorimotor boundary decisions, observer accounts, rival formulation	Small study	The reconstruction changes no diagnosis or prediction	Embodiment repairs the observer problem	4E cognition and organizational cybernetics collaborators	Embodied observer paper

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The model handles change of kind, not only change in environment	Specify a discriminating morphogenesis case against the current architecture, starting with Beer's Chilean redesign of the recursions[7]	Longitudinal transformation record and rival accounts	Small study / program	Existing System Four and recursion machinery explains the kind-change without residue	Morphogenesis is missing from the architecture	Development systems and organizational-change researchers	Kind-change study

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
The six Part Four residues are unique to the VSM	Adversarial comparison with named neighboring theories	Matched concepts, operational consequences, specialists from both sides	Program	A neighbor supplies the same work with equal or greater precision	The claimed remainder belongs uniquely to this tradition	Cross-theory panel with adversarial review	Six-residue comparison

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
Each of the five functions carries weight in viability that the others cannot supply	**Formative reanalysis of Michael D. Pfiffner's 135-case dataset**: respecify the measurement model, remove each function block in turn, compare against a rule published beforehand	Existing dataset (2017, N=135); the author's participation; a registered decision rule, including how variance shared between blocks is read	Desk / small study	A function block can be removed without loss of discrimination	The five functions are each necessary (tested in Beer's sense only where absence is coded as in Chapter 28)	Structural-equation specialist with the original author	New — added 14 September 2026

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
Viability requires the five functions	Seek deviant cases using an open coding protocol and configural analysis	Multi-case register; the coding in Chapter 28 (official, disseminated or absent, by registered indicators); viability outcomes independent of that coding, not bare persistence	Program	Viability kept with one function coded absent, or lost from an unseen cause with all five coded present	The configuration states a necessity of viable organization	Comparative organization and QCA researchers	Structural disconfirmation register

Claim at risk	Experiment	Evidence required	Minimum viable cost	Result that counts against us	Threatened assumption	Who could do it	Planned companion dossier
A single price per commodity lacks requisite variety for some regulatory tasks	Requisite-variety analysis of Hayek's price mechanism: does one scalar per commodity have requisite variety, and a short enough settling time, for each regulatory task (Chapter 16)	Classes of disturbance, a model of the price signal, settling-time evidence	Not yet priced	One scalar per commodity suffices in variety and settling time for every task examined	Requisite variety answers Hayek's objection	Economist with a control theorist	New — added 29 September 2026

The two necessity rows overlap. The formative re-analysis of Pfiffner's dataset is attempted first; the deviant-case search is what remains if the dataset cannot be obtained.

The table prices bands, not quotations. A band tells a reader which items can be refused on grounds of cost and which cannot, and it is fit for nothing more. It follows that no total has been struck. The receipts in the earlier chapters are stated in mixed units, and the bands here are a third unit again, so the components cannot be summed without inventing an exchange rate between them. We will not invent one. The bill is a list of components whose sum has not been computed, and printing a figure we cannot defend would be a worse fault than printing none.

Whether any of this happens will be decided in the column that names who could do it. Most of the competences listed there are not held inside this tradition. We cannot commission a neuroanatomist, a safety researcher, an information theorist or a comparative-methods team. We can only ask, and we should expect some to decline and others to accept and then produce results we would rather not have. That is the arrangement we are proposing. The cheap items at the top of the table are the ones we can start without anyone's permission, and starting there is the only evidence available that the offer is made in good faith.

The state of the count

The same discipline applies to what we currently know about our own exposure. Below are the figures we have, with their status stated before they are read.

As of August 2026

Both searches reported here are dated, and both will be wrong before this book is old. The arXiv full-corpus search returned three all-time hits for "viable system model" and none for either "management cybernetics" or "organizational cybernetics". *Kybernetes* had indexed four hundred and eighty-three papers since January 2024, of which none addressed neuroscience, predictive processing, active inference or allostasis.

A count of this kind is a photograph. What gives it value is the date on it, and the register is where the next photograph goes.

---

[3] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 132, where the sympathetic loop sits within System Three, against 176, "the sympathetic (System Two) circuits"; see also 129.

[4] Beer, *Heart of Enterprise*, 235 and 352, as quoted in Chapter 28.

[5] System Four must "contain a model of the viable system of which it is the System Four" (Beer, *Heart of Enterprise*, 235), raises "the issue of initiative" (237), provides "self-awareness" (Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 115), and runs a loop "anchored in THE UNKNOWN FUTURE" (119).

[6] Rate separation as a criterion of level is Herbert A. Simon's, in "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82; Chapter 11 discusses it and Chapter 29 sets out the design.

[7] Beer, *Brain of the Firm*, 322–26; on 323 he concludes that "two levels of recursion were mixed together in that preliminary model."

What we owe

The obligations that follow from the bill are narrower than the bill. We owe the work at the cheap end, because cost gives us no excuse for leaving it undone. We owe the three discriminators to anyone we ask to take an item, since without them the register becomes another literature nobody outside the tradition has reason to read. We owe the register an owner and a cadence, and it has both. And we owe the results that go against us the same treatment as those that do not. Of the four, that obligation is the easiest to meet in principle and the easiest to avoid in practice.

Every item in the bill is offered on those terms. Take one, try to break it, and report what broke, including to the people who would prefer it had held. The components are listed above; the total has not been struck; the register is where the account continues after this page.

Chapter 31

Coda: Judy

Judy survives because she gets up changed. That has been the argument of this book. We never claimed the model is sound. We claimed it is the kind of thing that could be corrected, and that we have not been correcting it. A boxer who gets up the same as she went down has learned nothing, and a model defended in the form it was inherited has been embalmed, however carefully it is kept.

One question is left, and it is about who her getting up is for.

A cartel can be viable

A cartel can be exemplarily viable. Clear identity; well-adapted operating units; a coordination layer that stops the members undercutting each other; an intelligence function watching for the antitrust authorities; a policy function that settles what the cartel is and will not become. A secret police can be viable. So can a firm whose product harms its customers slowly enough that they have not yet organized to say so. The model will help each of them survive, in the same vocabulary, with no change of tone. Beer's First Principle asks that varieties equate "with minimum damage to people and to cost", and his Establishment is a pathology of internal energy use; both protect those inside.[1]

Nobody constructed this objection to embarrass the model. It follows directly from what the model claims. Viability is a property of a system's capacity to persist under variety it cannot fully anticipate, and that capacity is indifferent to what the system is doing to anything outside itself. Nothing in the diagnosis distinguishes a hospital from a cartel. The recursion works the same way in both. The channels carry the same traffic. The audit function does the same job.

Neutrality would be bad enough, and the model is also *useful*. A theory that gave the cartel bad advice would limit itself: its own failures would keep it away from work we would not want it doing. This one does not fail there. Whatever you point it at, it makes better at persisting. That property

[1]Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 565, the First Principle of Organization as restated in the glossary (the statement on 97 reads "minimal damage"); and 412, where " 'The Establishment' in any social system comes into being at the point when the vital principal of autopoiesis consumes energy greater than that needed to maintain cohesiveness."

made the chapters before this one worth writing, and it makes this one unavoidable. An instrument that works can be used against people, and the better it works, the more that matters.

Viable for someone

A viable system is viable *for someone*. The model tells you whether the thing you pointed it at can keep going in a changing world. It does not tell you what to point it at. Naming the system, deciding that this is the thing whose survival is at stake, that these are its operating units, that its boundary runs here, is a decision taken before the modeling starts. Beer made it the method's first step: "determine precisely what is the organization to be modelled, and to specify its boundaries." He called a system "something defined by intelligence," wrote that two people "will not necessarily agree on the existence, or nature, or boundaries" of a system, and told the analyst not to "gaily accept the system delineated by the conventions of management structure."[2] Nothing in the model says whose interests that decision must answer to, and nothing records who made it.

Someone makes it, and that someone has interests. In practice it is whoever commissioned the work: the board, the ministry, the executive team, the client who is paying. The naming is the most consequential act in the whole engagement, and it is the one act the model names and leaves without a procedure. Which units get autonomy, which flows count as coordination, what the intelligence function is told to watch for, what the algedonic channel is wired to interrupt: all of it sits downstream of a choice the method treats as given.

Chapter 16, *The Objections That Land, and What They Cost*, sets out Werner Ulrich's twelve boundary questions, and Chapter 4, *Channels Without Owners*, sets three of them beside the channels the model does specify. The tenth is the one this chapter is left with: who is witness to the interests of those affected but not involved.[3] Beer names a place for an answer and never specifies it. "All stakeholders are represented in System Five," he wrote in 1994, and effectiveness is decided by "the perception and the purpose of the people who sponsor, operate, use and serve the organization."[4] Who counts as a stakeholder, and how the perception of those outside the membership enters, he does not say. The algedonic channel is the nearest

[2]Stafford Beer, *Diagnosing the System for Organizations* (John Wiley & Sons, 1985), 2 and 16; Stafford Beer, *Decision and Control* (John Wiley & Sons, 1966), 242, 243 and 246.

[3]The sources for critical systems heuristics are given in Chapter 16. The wording of the tenth question follows Ulrich's own checklist: "Who is (ought to be) witness to the interests of those affected but not involved?" (Werner Ulrich, "A Brief Introduction to Critical Systems Heuristics (CSH)," 14 October 2005, fig. 4, https://wulrich.com/downloads/ulrich_2005f.pdf).

[4]Stafford Beer, *Beyond Dispute: The Invention of Team Syntegrity* (John Wiley & Sons, 1994), 149; Stafford Beer, *Think Before You Think*, ed. David Whittaker (Wavestone Press, 2009), 24. See also Stafford Beer, *The Heart of Enterprise* (John Wiley & Sons, 1979), 263: "In a nationalized industry, a 'consumer's council' may embody the role of System Five."

thing it has to a mechanism for a signal that overrides the ordinary hierarchy of attention, and it has no seat for the person who is affected and was never a member.

Different logical levels

Nothing in critical systems heuristics is incompatible with the Viable System Model. The two operate at different logical levels, and the long habit of treating them as rivals[5] has cost this tradition a repair proposed decades ago. Jackson proposed this division of labor in 1988, with critical systems heuristics "a better guide" to power, and Flood and Jackson built it into Total Systems Intervention in 1991; the tradition did not take it up.[6] Ulrich interrogates the boundary. Beer designs inside it. A design method that names its boundary as an observer's act and gives no procedure for examining it, and a boundary critique with no design apparatus: each is the other's missing half.

The slot is nearly there. System Five and the algedonic channel are, in our reading, a witness function, the point at which something that matters enough to override the routine can reach the level able to act on it. In the firm, though, Beer designs the channel as a statistical alarm on points of stability that the metasystem chooses.[7] Only the Chilean meter let persons trip it, and those persons were members of the polity it served.[8] The model has the mechanism and lacks a theory of who is entitled to use it. Membership is assumed. The signal is expected to originate somewhere inside, from a unit whose distress is the system's distress. There is no account of standing for a party whose distress is caused by the system operating exactly as designed.

The defect is fixable, and it does not refute the model. It is also the cheapest high-value reform available to this tradition. It needs no new formalism, only a statement of who has standing to trip the channel and who sets what it watches, built into the diagnosis so that its absence shows as an absence instead of passing unnoticed.

---

[5] M. C. Jackson, "An Appreciation of Stafford Beer's 'Viable System' Viewpoint on Managerial Practice," *Journal of Management Studies* 25, no. 6 (1988): 557–73, records the dispute: Ulrich 1981 and 1983 among the "severe criticism" of the model (557), Beer's 1983 reply to Ulrich (564), and Ulrich's charge that Beer practiced "tool design" rather than "social system design" (568–69).

[6] Jackson, "An Appreciation," 571: the VSM contributes most to the technical interest, and on power "Ulrich's (1983) 'critical systems heuristics' is a better guide here." Robert L. Flood and Michael C. Jackson, *Creative Problem Solving: Total Systems Intervention* (John Wiley & Sons, 1991), as described in Michael C. Jackson, *Critical Systems Thinking and the Management of Complexity* (John Wiley & Sons, 2019), 540–43, where methodologies are chosen and combined according to the problem context; we have not read the 1991 book itself.

[7] Stafford Beer, *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981), 178–79; Beer, *Heart of Enterprise*, 275 and 381–86; Beer, *Diagnosing the System*, 133.

[8] Beer, *Brain of the Firm*, 283–84.

The standing test

Test it yourself on the last thing you built or approved. Leave aside "did we engage stakeholders" and ask: if every person affected by this had said nothing at all, what would have happened? If it would have gone ahead exactly as it did, then whatever else was there, standing was not.

The test is deliberately crude. Consultation that changes nothing does not amount to standing. It is a channel with no variety-absorbing capacity, which this tradition already has the language to describe and has rarely turned on its own practice. Standing means a signal from outside the membership can change the outcome. Where it cannot, the channel is decorative.

What the book has been walking toward

Viability is not a value. It is a capacity, and the model makes whatever it is pointed at better at holding it. The choice of what to point it at is made outside the model, by someone with interests, and the model names that choice without examining it or leaving a record of it.

That is the reason the repair set out in Chapter 4 matters, and the reason the agenda in Part Five matters. An instrument this useful, in a tradition that has never stated what would count against it, is a practical hazard. A community that cannot say what observation would make it give something up cannot say what observation would make it decline a piece of work either. Both need the same thing: a statement, fixed in advance, of what the model does not cover and whom it does not serve. The agenda of scrutiny and the question of standing are one discipline, applied at two points.

This is why the allegory ends where it does. Judy stays in the fight because she is corrigible, and corrigibility belongs to a practice, not to a doctrine. Her getting up again does not tell us whether she should stay in the fight. Persistence answers whether she can. It does not answer who she is fighting for, and that question is not asked inside the ring.

So the book's last ask is small, and it is addressed to those of us who use this model in engagements.

If you are in this tradition: ask the question in one engagement, in writing. Whose residual variety is this intervention expected to increase, and whose will it reduce? So far as we have been able to establish, and without the systematic survey of the method literature that Chapter 26 says has not been done, no method book in this field asks it in this form: not the ones we teach from, and not the ones we wrote. Beer supplied its terms. Freedom, in *The Heart of Enterprise*, lies "in what the institution DOES to them by way of variety constraint," and *Diagnosing the System* asks how the First Axiom looks "as System One sees it."[9] *Residual variety* here has the sense defined in Chapter 26, *Freedom as a Quantity*: the room left to a person

[9] Beer, *Heart of Enterprise*, 156–58 and 160; Beer, *Diagnosing the System*, 95. Chapter 26 notes the exercise on freedom at *Diagnosing the System*, 34.

or a unit, not Beer's Third Axiom residue. The machinery for answering the question is already there. Cost: one client may answer honestly, and then you know.

Glossary

Terms marked with ↗ in the text are defined here on first appearance in each chapter.

Editions and pagination. Beer is cited by printed page in these editions: *Brain of the Firm*, 2nd ed. (John Wiley & Sons, 1981); *The Heart of Enterprise* (John Wiley & Sons, 1979); *Diagnosing the System for Organizations* (John Wiley & Sons, 1985); *Designing Freedom* (Canadian Broadcasting Corporation, 1974), cited by lecture because the copy used is unpaginated. Other authors are cited as in the chapter notes. Each entry gives Beer's usage first and the book's usage only where the book deliberately differs or narrows it.

A

Active inference — The extension of Karl Friston's free energy principle to action: a system that maintains itself either changes its beliefs to match the world or changes the world to match its beliefs. The book takes it seriously because it runs: there is working code and a pipeline from generative model through simulation to fit against behavioral data (Ch 12). Friston, FitzGerald, Rigoli, Schwartenbeck and Pezzulo, "Active Inference: A Process Theory," *Neural Computation* 29, no. 1 (2017): 1–49; Parr, Pezzulo and Friston, *Active Inference* (MIT Press, 2022). See also *Free energy principle*. See Ch 10, 12, 21, 27.

Adaptive cycle — see *Panarchy*.

Algedonic signal, algedonic channel — Beer's term (Greek *algos*, pain, and *hedos*, pleasure) for regulation "in a non-analytical mode" that raises alarm. Algedonic signals, "alone of all the signals in the enterprise," break through the levels of recursion; a filter that "monitors chosen points of stability by statistical methods," with criteria set by the metasystem, decides what reaches System Five. *Diagnosing the System for Organizations* (1985), 134; *The Heart of Enterprise* (1979), 386, 407; *Brain of the Firm* (1981), 178–79. The book's point is that the channel is specified as a mechanism and never as an institution: no one is named who chooses the points of stability, contests the filter or audits it (Ch 4); Ulrich reads the same signal as a transfer of control upward (Ch 16). See Ch 4, 8, 9, 11, 16, 24.

Amplifier, variety amplification — In Beer, "a device that increases variety"; with the attenuator, one of the two moves by which varieties are balanced across a boundary. *Diagnosing the System for Organizations* (1985),

35. Galbraith's "increase the capacity to process information" is the same move reached without the vocabulary (Ch 12). See Ch 11, 12.

Ashby's Law — see *Requisite variety, law of*.

Attenuator, attenuation — In Beer, "a device that reduces variety." Every regulator faces more variety than it can hold and must attenuate; Beer called the recursive description of an organization itself "a vast attenuator of the variety proliferated by any large organization." *Diagnosing the System for Organizations* (1985), 35; *The Heart of Enterprise* (1979), 308. The book keeps the technical sense and, following *Designing Freedom*, stresses that attenuation is something done to people; what is left after it is their residual variety (Ch 26). See Ch 11, 12, 13, 16, 26.

Audit — see *System Three-Star*.

Autonomy — In Beer, "the freedom of an embedded subsystem to act on its own initiative, but only within the framework of action determined by the purpose of the total system." The degree of autonomy is agreed in the resource bargain. Beer held that "freedom is in principle a computable function of systemic purpose as perceived," and gave no computation. *Diagnosing the System for Organizations* (1985), 38, 105; *The Heart of Enterprise* (1979), 158, restated at 200. See *Autonomy–cohesion problem*. See Ch 10, 13, 16, 26.

Autonomy–cohesion problem — The problem of setting how much freedom the parts of a system keep: too much and they drift out of alignment, too little and the center must absorb variety it cannot hold. Beer's rule is that the metasystem "should make only that degree of intervention that is required to maintain cohesiveness." *The Heart of Enterprise* (1979), 158. The book finds the same tension in Koestler's holon and working mechanisms for it in sociotechnical design and sociocracy, and records that nobody has supplied the computation, measure or thresholds the rule needs (Ch 11, 13). See Ch 9, 11, 13, 24.

B

Boundary judgment — In Werner Ulrich's critical systems heuristics, the judgment, always made by someone, of what belongs inside a system under study and who is affected by it without being consulted. Ulrich, *Critical Heuristics of Social Planning* (Paul Haupt, 1983). The book uses it to ask who is entitled to name a system's identity, purpose and levels (Ch 16, Coda). See Ch 4, 10, 16, Coda.

C

Closure — In Beer, "a self-referential process, and not the isolation of the system within an adiabatic shell"; closure "turns the system back into itself," and "Closure is the talisman of identity." System Five supplies "logical closure to the viable system." *The Heart of Enterprise* (1979), 260; *Diagnosing the System for Organizations* (1985), 129. See Ch 10, 16.

Cohesion, Law of Cohesion — The holding together of a viable system and of its levels of recursion. Beer's Law of Cohesion for Multiple Recursions: "The System One variety accessible to System Three of Recursion x = the variety disposed by the sum of the metasystems of Recursion y." He called planning the "glue" of organizational cohesion and held that cohesion "governs the question of identity in the viable system." *The Heart of Enterprise* (1979), 336, 351, 355; *Diagnosing the System for Organizations* (1985), 134. See *Autonomy–cohesion problem*, *Six vertical components*. See Ch 4, 10, 13, 16.

D

Decerebrate cat — Beer's image for an institution that goes on existing without a working metasystem: "It lives on; its viability is ensured by a bogus environment, and is sustained by artificial sustenance." *The Heart of Enterprise* (1979), 265. The book uses it to separate Beer's viability from bare persistence (Ch 9, 28). See Ch 9, 15, 18, 28.

Double link — see *Sociocracy*.

F

Free energy principle — Karl Friston's claim that any system that persists, and whose states divide by a Markov blanket into internal, sensory, active and external, will appear on average to minimize variational free energy, so that its internal dynamics can be read as approximate Bayesian inference about its environment. Friston, "The Free-Energy Principle: A Unified Brain Theory?," *Nature Reviews Neuroscience* 11, no. 2 (2010): 127–38; Friston, "Life as We Know It," *Journal of the Royal Society Interface* 10, no. 86 (2013): 20130475. The technical criticism of its formal argument is set out in Ch 12. See Ch 10, 12, 21.

Freedom — see *Autonomy*; *Residual variety*.

G

Good regulator theorem — Conant and Ashby's result, stated in their title: "Every good regulator of a system must be a model of that system." Its measure of success is minimal entropy of outcomes, which a reliably bad regulator also meets, and the authors conceded that some optimal regulators are not models of what they regulate. Conant and Ashby, *International Journal of Systems Science* 1, no. 2 (1970): 89–97, at 91 and 96. See Ch 3, 12.

H

Holon — Arthur Koestler's term for any part that is also a whole. Every holon has "the dual tendency to preserve and assert its individuality as a quasi-autonomous whole; and to function as an integrated part of an…

larger whole." Koestler, *The Ghost in the Machine* (Hutchinson, 1967), where the term is coined (48); the quoted proposition is 4.1 of the book's Appendix I (343), given here in the wording of the 1969 Alpbach appendix, "Some General Properties of Self-Regulating Open Hierarchic Order," in *Beyond Reductionism*, ed. Koestler and Smythies (Hutchinson, 1969). A holon need not be viable, which separates it from Beer's recursion (Ch 10, 11); holonic manufacturing took the term into engineering (Ch 5, 15). See Ch 5, 10, 11, 15.

Homeostasis, homeostat — In Beer, homeostasis is the "stability of a system's internal environment despite the system's having to cope with an unpredictable external environment." In a homeostat two systems are coupled so that "each system acts (having of course requisite variety) as controller of the other." The Three–Four homeostat couples the inside-and-now (System Three) with the outside-and-then (System Four) and is monitored by System Five. *Diagnosing the System for Organizations* (1985), 17, 111, 129; *Brain of the Firm* (1981), 145. See Ch 2, 8, 16.

I

Identity — What makes a viable system the same system through change. Beer ties it to closure and to cohesion, and treats boundaries and purposes as observer-dependent: "The facts about the system are in the eye of the beholder." *The Heart of Enterprise* (1979), 9, 260, 351. The book notes that the model does not say who is entitled to name a system's identity or how a disagreement about it is settled (Ch 16, Coda). See Ch 7, 11, 16.

L

Law of requisite hierarchy — Aulin-Ahmavaara's result that hierarchy compensates when regulatory ability is limited. The hierarchy it derives is one of control, and the paper mentions neither Beer nor recursion. Aulin-Ahmavaara, "The Law of Requisite Hierarchy," *Kybernetes* 8, no. 4 (1979): 259–66. See Ch 3, 12.

Level of recursion — see *Recursion*.

Liberty machine — In *Designing Freedom*, government redesigned "to be, not an entity characterized by more or less constraint, but a dynamic viable system that has liberty as its output." *Designing Freedom* (1974), lecture 3, "A Liberty Machine in Prototype." The book reads its output as the residual variety of the people inside it (Ch 26). See Ch 26.

M

Markov blanket — Judea Pearl's construct, an epistemic tool for inference in a network; in the free energy literature, the partition that separates a system's internal states from external ones. Its route into active inference runs through Bayesian networks and variational inference, not through

cybernetics, and Pearl's blanket and Friston's metaphysical use of it are different constructs. Pearl, *Probabilistic Reasoning in Intelligent Systems* (Morgan Kaufmann, 1988); Bruineberg, Dołęga, Dewhurst and Baltieri, "The Emperor's New Markov Blankets," *Behavioral and Brain Sciences* 45 (2022): e183. See Ch 12, 21.

Metasystem — In Beer, "a system 'over and beyond' a system of lower logical order. (Higher authority is not the issue, and may not apply.)" In the VSM, Systems Three, Four and Five form a structure "metasystemic to the 3-2-1." *Diagnosing the System for Organizations* (1985), 129, 134. The book notes that Beer's metasystem is not optional, unlike the staff holon of holonic manufacturing (Ch 5, 15), and that it works by attenuating (Ch 16). See Ch 4, 5, 13, 15, 16.

Minimum critical specification — The sociotechnical design rule that a specification should state no more than is essential and leave the rest to those who do the work. The principle is P. G. Herbst's; Albert Cherns named and systematized it while disclaiming originality. Herbst, *Socio-Technical Design* (Tavistock Publications, 1974); Cherns, "The Principles of Sociotechnical Design," *Human Relations* 29, no. 8 (1976): 783–92. The book reads it as a procedure for System Three's job, still without a measure of how much specification is critical (Ch 13). See Ch 10, 13.

N

Near-decomposability — Herbert Simon's definition of hierarchy: the parts within a subsystem interact much faster or more strongly than subsystems interact with each other. Simon explicitly includes systems with no relation of subordination, and the threshold that marks a level is chosen by the analyst. Simon, "The Architecture of Complexity," *Proceedings of the American Philosophical Society* 106, no. 6 (1962): 467–82, at 468. See Ch 10, 11, 29.

O

Organizational information-processing theory (OIPT) — Jay Galbraith's account in which structure is set by the information that must be processed under uncertainty, so that a designer can either reduce the need for information or increase the capacity to process it. Galbraith, *Designing Complex Organizations* (Addison-Wesley, 1973), and "Organization Design: An Information Processing View," *Interfaces* 4, no. 3 (1974): 28–36; operationalized with validated constructs from Tushman and Nadler, *Academy of Management Review* 3, no. 3 (1978), onward. The book reads it as variety attenuation and amplification reached without Ashby or Beer (Ch 10, 12). See Ch 10, 12, 29.

Oscillation — In Beer, "failing to settle down in homeostatic equilibrium, a dynamic system over-corrects itself continuously." Damping it is the

function of System Two. *Diagnosing the System for Organizations* (1985), 71. See Ch 13, 24, 28.

Ostrom's design principles — Elinor Ostrom's eight principles for long-enduring common-pool resource institutions: clearly defined boundaries; congruence of rules with local conditions; collective-choice arrangements; monitors accountable to the appropriators; graduated sanctions; low-cost local conflict resolution; minimal recognition of the right to organize; and, for resources that are parts of larger systems, nested enterprises. Ostrom, *Governing the Commons* (Cambridge University Press, 1990), 90 (table 3.1). The eighth resembles recursion without requiring each layer to be viable; the third is a legitimacy condition on System Five that Beer did not impose (Ch 14). See Ch 4, 5, 10, 14.

P

Panarchy — The account by C. S. Holling and Lance Gunderson of adaptive cycles (exploitation, conservation, release, reorganization) nested across scales and coupled in two directions: *revolt*, a fast collapse at a small scale cascading into a slower, larger level, and *remember*, the slow large level supplying structure to rebuild a small one. Gunderson and Holling, eds., *Panarchy* (Island Press, 2002); Holling, "Understanding the Complexity of Economic, Ecological, and Social Systems," *Ecosystems* 4, no. 5 (2001): 390–405. It comes after Beer, so it is convergence rather than priority; Beer's couplings between levels do not depend on phase, and nothing in the model corresponds to *remember* (Ch 10, 11). See Ch 4, 5, 10, 11, 27.

R

Recursion, level of recursion — Beer's Recursive System Theorem: "In a recursive organizational structure, any viable system contains, and is contained in, a viable system." Recursion is a relation between viable systems, not another word for organizational level: "If the term *recursion* is allowed to become a meaningless appendage to the term *organizational level*, then we do not have a valid model of a viable system"; and "Systems Two, Three, Four, and Five are not in themselves viable systems at any level of recursion." *The Heart of Enterprise* (1979), 118, 308, 411; *Diagnosing the System for Organizations* (1985), 2, 17. The book keeps this criterion, finds no procedure that fixes where the levels lie or who fixes them (Ch 5, 16), and finds that neighboring literatures reached nesting, not recursion (Ch 10, 11, 14). See Ch 5, 10, 11, 13, 14, 16, 24, 29.

Requisite variety, law of (Ashby's Law) — W. Ross Ashby's law that only variety in a regulator can reduce the variety of outcomes: "only variety in R can force down the variety due to D; variety can destroy variety." Ashby, *An Introduction to Cybernetics* (Chapman & Hall, 1956), 207. Beer's compression, "only variety can absorb variety," turns Ashby's bound into a design instruction. *Diagnosing the System for Organizations* (1985), 35. The book keeps the

two apart and records the conditions on which Ashby's derivation depends (Ch 3). See Ch 3, 12, 16, 26.

Residual variety — Two senses, which the book keeps apart. (1) Beer's: the variety left over from the balance of Systems Three and Four, which System Five absorbs. His Third Axiom of Management reads: "The variety disposed by System Five is equivalent to the residual variety generated by the operation of the Second Axiom." *The Heart of Enterprise* (1979), 216, 263, 298, restated with "equals" at 567; *Diagnosing the System for Organizations* (1985), 130. (2) The book's, in Part Four: the variety left to a person or an operational unit after an institution's attenuation has taken what the institution needs to stay stable. *Designing Freedom* describes this quantity without naming it; in *The Heart of Enterprise* Beer called it freedom, "a computable function of the purposes of the system as perceived" (200). See Ch 26, 27, 28, Coda.

Resource bargain — In Beer, "the 'deal' by which some degree of autonomy is agreed between the Senior Management and its junior counterparts"; one of the six vertical components. *Diagnosing the System for Organizations* (1985), 38. See Ch 4, 13, 16.

S

Six components — The book's grouping of the design ideas that neighboring literatures reached independently of Beer: variety matching (Galbraith), nesting (Simon, Koestler), criteria for levels (Jaques), cross-scale coupling (Holling), bounded autonomy (Trist and the sociotechnical tradition) and nested governance (Ostrom). The grouping is the authors', not Beer's six vertical components. See Ch 10, 13, 14, 23.

Six vertical components — Beer's six vertical channels through which the metasystem absorbs the variety generated by the operational elements: corporate intervention (command), the resource bargain, operational linkages, environmental intersects, System Two and audit. His First Axiom of Management balances horizontal variety against "the six vertical components of corporate cohesion." *The Heart of Enterprise* (1979), 143, 217; *Diagnosing the System for Organizations* (1985), 64, 82, 84. See Ch 4, 10, 13.

Sociocracy, double link — Gerard Endenburg's circle method, built on Kees Boeke's consent principle: a decision stands when no member raises an argued objection; work is done in semi-autonomous circles; adjacent circles are joined by two people, a leader appointed downward and a delegate elected upward (the double link); and people are assigned to functions by consent. Endenburg, *Sociocracy* (Eburon, 1998); Boeke, "Sociocracy: Democracy as It Might Be" (1945); Romme and Endenburg, *Organization Science* 17, no. 2 (2006): 287–97. The book reads the double link as a variety design and a legitimacy design at once, tested on neither count (Ch 13). See Ch 10, 13.

System-in-focus — Beer's term for the viable system chosen for study, "central to a whole sphere of existence" marked out by its higher and lower

recursions, of which it may have more than one. *Diagnosing the System for Organizations* (1985), 6. See Ch 5, 16.

System One — The operational elements of a viable system, the units that do what the system exists to do; each is itself viable. "System One consists of viable systems whose conditions of survival are formulated at a different level of recursion." *Diagnosing the System for Organizations* (1985), 100; *The Heart of Enterprise* (1979), 411. The book notes that what counts as an instance of System One at a given level is settled case by case (Ch 5). See Ch 5, 10, 11, 13, 16.

System Two — "the viable system's anti-oscillatory device for System One." "Its function is not to command, but to damp oscillations," and it "is always present, or the organization would shake itself to pieces." *Diagnosing the System for Organizations* (1985), 66, 68. See Ch 1, 24, 28.

System Three — The management of the "inside-and-now": it oversees System One and System Two as a whole, strikes the resource bargain, and should intervene only as far as cohesion requires. *Diagnosing the System for Organizations* (1985), 38, 111; *The Heart of Enterprise* (1979), 158. Its balance with System Four is the Three–Four homeostat. See Ch 5, 11, 13, 16.

System Three-Star (3*), audit — In *Diagnosing the System*, the sporadic, high-variety channel by which System Three reaches past the managers of System One and, "by agreement with System One management," penetrates to the operations; audit is the typical case. *Diagnosing the System for Organizations* (1985), 82, 86. In *The Heart of Enterprise* the label 3* names something else, a model of System Three held inside System Five (352). See Ch 1, 4, 14.

System Four — The management of the "outside-and-then," which "houses the viable system's whole apparatus for adaptation" and "must contain a model of the viable system of which it is the System Four." "System Four exists — it has to: but it is disseminated, and lacks integration." *The Heart of Enterprise* (1979), 143, 235; *Brain of the Firm* (1981), 183; *Diagnosing the System for Organizations* (1985), 111, 115. The book argues that teaching has reduced it to a pair of eyes scanning the environment, whereas Beer's System Four models and anticipates (Ch 2). See Ch 2, 3, 16, 28.

System Five — The function that supplies "logical closure to the viable system" and monitors the Three–Four homeostat; "the metasystemic administrator of Ashby's Law," which "necessarily absorbs the residual variety of the Three-Four interaction." "Under more democratic conventions, System Five consists of many people." *Diagnosing the System for Organizations* (1985), 129; *The Heart of Enterprise* (1979), 263. See Ch 4, 6, 14, 16, 26.

T

Three–Four homeostat — see *Homeostasis, homeostat*.

Time-span of discretion — Elliott Jaques's measure of the level of work in a role: the longest period over which a person exercises discretion before its results become reviewable. Jaques held that strata are discrete, with

boundaries at about three months, one year, two, five, ten and twenty years. Jaques, *Measurement of Responsibility* (Tavistock Publications, 1956); *Requisite Organization* (Cason Hall, 1989). Its strata are layers of managerial roles, the reading of level that Beer ruled out for recursion (Ch 10, 11). See Ch 5, 10, 11, 29.

Transducer — In Beer, a device that "encodes or decodes a message whenever it crosses a system boundary." His Third Principle of Organization: "Wherever the information carried on a channel capable of distinguishing a given variety crosses a boundary, it undergoes transduction; and the variety of the transducer must be at least equivalent to the variety of the channel." Channel capacity and transduction capacity "are not at all the same thing." *The Heart of Enterprise* (1979), 101; *Diagnosing the System for Organizations* (1985), 47, 53. The book finds the principle stated and never operationalized or measured (Ch 25). See Ch 3, 4, 25.

V

Variety — Ashby's measure, used by Beer: "a measure of complexity: the number of possible states of a system." *Diagnosing the System for Organizations* (1985), 35. Ashby added that a set's variety depends on the observer's powers of discrimination (*An Introduction to Cybernetics*, 125), and Beer conceded that the variety of a real organization cannot be counted: the use "is not the naive use of counting states, but of matching state generators." *The Heart of Enterprise* (1979), 86–87. See Ch 3, 12, 13, 25, 26.

Viable, viability, viable system — Beer takes "viable" from the *Oxford English Dictionary*: "able to maintain a separate existence." *Diagnosing the System for Organizations* (1985), 1, 17. The outcome in the book's proposed tests is Beer's viability, not bare persistence, since a system can persist without a working metasystem (see *Decerebrate cat*; Ch 28). See Ch 5, 9, 10, 18, 28.

Viable System Model (VSM) — Beer's model of the organization of any viable system: five subsystems, System One to System Five, repeated at every level of recursion. Beer called the choice of five "somewhat arbitrary" and held that "a model is neither true nor false: it is more or less useful." *Brain of the Firm* (1981), 98; *Diagnosing the System for Organizations* (1985), 2. The book treats the VSM as an assembly of components that other literatures reached independently, and narrows its claim accordingly (Ch 10, 23). See Ch 10, 12, 13, 16, 23.

Acknowledgements

About the Authors

Krishan Mathis is a consultant, author, and practitioner-scholar working at the intersection of management cybernetics, systems thinking, and organisational design. Through his practice, tautai, he helps organisations grow strategic adaptivity — navigating complexity without relying on prediction. He is the author of *The Tautai Principle* (2025) and *Education is Broken* (2026), on human work in times of AI.

Margeret Heath is an independent researcher and self-described generalist whose inquiries centre on relational and interactive systems, cybernetics, and niche-construction processes. Her work investigates how dynamic, nested systems self-regulate and generate novel affordances amid uncertainty, with particular interest in distributed embodied social cognition and the collective playful imaginary. She pursues transdisciplinary collaborations toward regenerative futures and planetary flourishing.