Who holds the Pen

Krishan Mathis

Who Holds the Pen

Krishan Mathis

Tautai Press

First edition — 2026-05-05

Tautai Press
First edition — 2026-05-05

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

Contents

Reading the Water

An Introduction — Why One Model Needs Seven Methods, and How to Choose Among Them

I wrote *The Tautai Principle* because I had come to believe that adaptive organizations are navigated, not planned — that in a world which rewards sensing and response over prediction and control, the organizations that win are the ones that feel a shift in their environment and move on it weeks before their rivals. The book's guiding image is the Tautai, the Polynesian wayfinders who crossed the largest ocean on earth without charts, reading birds and swells and the colour of the water rather than following a fixed course.

One chapter of that book leans directly on a piece of management cybernetics I have never stopped finding useful. In Chapter 5, "The Organizational Nervous System," I use Stafford Beer's Viable System Model as the blueprint for designing the pathways along which intelligence has to travel if it is going to move at market speed rather than hierarchy speed. The VSM is a rare thing: a genuine theory of what makes an organization *viable* — capable, in Beer's sense, of adapting and surviving in a changing environment. It gives you five functions every viable system must perform, the channels that connect them, and the same pattern repeating at every level of scale. I called it the most powerful blueprint available for the job, and I still would.

But I was careful, even there, to name what the model does not do. Its focus is formal structure. It will show you that a signal is being lost between the front line and the people who could act on it, and it will tell you which channel is broken — but it is quiet on the human factors, the culture, the judgement that decides whether anyone acts at all. The VSM builds the resilient ship; it does not sail it. That is why, in the book, I pair it with other tools. And it is why, when I later set out to actually *use* the model on a real organization — not to explain it, but to run a diagnosis or design an intervention — I ran into a gap the model itself never closes.

The Model Is an Anatomy, Not a Procedure

Here is the gap, stated plainly. The VSM tells you, with unusual precision, what a viable system *looks like*. It does not tell you how to *get there*. It is an anatomy, not a procedure. It shows you the destination — the five systems

in balance, the channels carrying the right variety, the recursion clean at every level — and leaves the navigation to you.

So I went looking for the method: the step-by-step, the "what do I do on Monday morning." What I found was not a method. It was a shelf of them.

And the surprise — the thing that turned a reading list into the set of chapters that follows — was that the authors of those books disagree far less about the model than I expected. On the model itself they are almost in unison: five systems, requisite variety, recursion, the algedonic alert. Where they part company is on a question the model never asks and Beer's own handbook never quite settles: *who holds the pen?* Who actually does the work of the analysis? One manager alone at a wall of hand-drawn charts? A small team designing in confidence before anyone else is told? An expert analyst with a piece of software? A skilled facilitator convening the people who cannot agree what the organization even is? Or the whole organization, in workshops, drawing its own map of itself?

Each of those is a different answer, and — this is the part that matters — each is a different *politics*. The choice of who holds the pen is a choice about where authority over the organization's self-description sits, and about how many people have to agree before the analysis can become the organization's reality. Three gaps, then, in the bare model: it gives no procedure, it supplies no judgement, and it is silent on this politics of authorship. The seven methods all close the first gap — that is what makes them methods. What separates them, and what this book is really about, is how each one answers the third.

Seven Instruments for Reading the Water

This set profiles seven serious methods for putting the VSM to work, each in its own chapter, each held to the same questions: what exactly is the method, what is genuinely unique about it, where is it at its best, and — the question I care about most — what does it assume about power, both in the organization and in the world the organization sits in?

They run, roughly, from the leanest to the widest constellation of people. Beer's own *Diagnosing the System* is the do-it-yourself handbook, written for one manager working the charts alone. Pfiffner's *Neurology of Business* hands the pen to a confidential core team of three to five, measuring where autonomy should sit instead of arguing about it. Hoverstadt's *Fractal Organization* is the consultant's field guide, diagnosing an organization by the way it fails — a bestiary of named pathologies you can recognize before you have learned a line of theory. Pérez Ríos's *Design and Diagnosis for Sustainable Organizations* is built around an instrument, a piece of software that makes a large, multi-level analysis tractable. Espejo and Reyes's VIPLAN begins by insisting that naming what the organization *is* — and for whom — is the first and most consequential act of the whole exercise. Schwaninger's *Intelligent Organizations* sets the VSM inside a larger frame, asking not only how variety is distributed but what a leadership team should actually be

steering by. And Espinosa and Walker's self-transformation method gives the pen to the whole community, so that the organization produces its own diagnosis and keeps the language after the consultant leaves.

A word of warning I wish someone had given me: four of these books carry "sustainable" or "sustainability" in or near their titles, and only *one* of them — Espinosa and Walker — uses the word to mean ecology. For the others it is a synonym for *viable*, a structural property of an organization that keeps its identity through change. It is the most common misreading in the literature, and it is worth clearing up before you choose.

When the Pen Is a Machine

For most of these methods' lives, the question "who holds the pen?" was almost rhetorical, because only a person could hold it. Running a VSM analysis took a scarce, trained human — Beer's manager at the wall, Pfiffner's core team, Pérez Ríos's analyst. That is no longer true: a capable AI model will now draft the whole diagnosis in minutes, and the pen can pass to a machine quietly, by default, without anyone deciding that it should. What that does to Beer's case for autonomy, to the participative methods, and to the one thing you least want to delegate — identity — is the argument of the coda that closes this book, and the practical question of Showcase 6. Here it is enough to say: the old question has stopped being rhetorical.

How to Read This Set

You do not have to read the seven in order, and you certainly do not have to adopt all of them. One convention throughout: bare chapter numbers refer to this book; references to *The Tautai Principle* always name it. Part II is the decision: a capstone that lays the seven side by side across the axes that separate them and turns the comparison into a decision guide — given a real organization and a real problem, which instrument should you reach for first. If you want the shortest possible route in, start there and let it send you to the two or three method chapters that fit your situation.

Part III — *The Model at Work* — then changes register. It collects six showcases: worked scenes in which a single VSM mechanism does its job — the balance between running today and adapting for tomorrow; a market signal that dies on its way up the hierarchy and the channel that saves it; a resource bargain that turns autonomy into accountability; the designed conversation between one recursion level and the next; narrative as the channel that steers thousands of decisions without instructions; and the placing of AI across the five systems. Their shared move — one this whole book keeps returning to — is to stop trying to declare a winner between two poles and to design the exchange between them instead. They are there to show the machinery running, not just described. Read them after the methods, or dip into whichever mechanism you are wrestling with.

The book closes with a coda, *Designing Freedom, Again*, which steps back to what the whole set is really about: freedom, and the distribution of agency in the age of AI. It is the destination the pen question has been walking toward, and it is best read last — after you have seen the machinery run.

I have kept one bias throughout, and I should name it. I read these methods as a navigator reads the water. A wayfinder is not handed a map; the whole art is to read what is actually in front of you and choose the next move from that. These seven methods are not seven maps. They are seven instruments for reading the same organizational water — and the discipline this set is really trying to teach is not any one of them, but the judgement to pick the right instrument for your vaka, your crew, and the sea you are actually on.

That judgement begins with a single, honest question, asked before any chart is drawn: *who holds the pen — and who, in the end, has to agree with what it writes?*

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1972). *Brain of the Firm*. London: Allen Lane. 2nd ed., Chichester: Wiley, 1981.
- Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.

- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5.

Part I

Part I — The Seven Methods

Seven Ways to Hold the Pen

The introduction ended on a question: who holds the pen, and who, in the end, has to agree with what it writes? Part I is seven answers to it. Each chapter takes one book that puts Stafford Beer's Viable System Model to work and profiles the method it proposes — not as a summary, but as a working assessment a practitioner could act on.

The chapters follow a common template, so they can be read against each other. Each opens with a real, worked case from the book itself; sets out the method step by step; names what is genuinely unique about it; and states its best application case and its limits. Then each does something the books themselves mostly decline to do: it makes the politics explicit — the external environment the method quietly assumes, and the internal power constellation it presupposes — before closing with a short "what to borrow." The recurring question underneath all of it is the one the set is named for.

The seven are arranged along a single spectrum: not from worst to best, but from the leanest pen-holder to the widest. One manager alone at a wall of charts, at one end; the whole community drawing its own map, at the other. Reading them in that order is the fastest way to feel the real variable — how many people must agree before an analysis can become the organization's reality.

- **Beer —** *Diagnosing the System for Organizations* (ch. 1): the originator's do-it-yourself handbook; one manager works the four Charts. The canonical source method.
- **Pfiffner —** *The Neurology of Business* (ch. 2): a seven-step process that measures variety instead of arguing; a confidential core team of three to five designs, then a reserved decision.
- **Espinosa & Walker —** *A Complexity Approach to Sustainability* (ch. 3): the self-transformation method; the organization co-constructs the model over months of workshops.
- **Pérez Ríos —** *Design and Diagnosis for Sustainable Organizations* (ch. 4): the most procedurally explicit method; an analyst works a four-stage sequence, a critical-factors matrix and a pathology catalogue, carried by VSMod® software.
- **Espejo & Reyes —** *Organizational Systems* (ch. 5): the VIPLAN method; identity-first diagnosis (TASCOI), the unfolding of complexity, and the recursion/functions table.
- **Schwaninger —** *Intelligent Organizations* (ch. 6): the VSM integrated with the Model of Systemic Control and its three logical levels of management.
- **Hoverstadt —** *The Fractal Organization* (ch. 7): the practitioner's field guide; diagnosis by symptom through a catalogue of structural pathologies.

You do not have to read them in order, or all of them. If you would rather start from your own problem than from the shelf, skip ahead to the capstone in Part II — it lays the seven side by side and sends you to the two or three that fit the water you are actually on.

Each chapter carries its own worked case, references, and companion chapters in The Tautai Principle. *Part II chooses among these seven; Part III shows the mechanisms they diagnose running in practice; the coda that closes the book asks what the choice is really about.*

Chapter 1

From the Horse's Mouth

Beer's Own Diagnostic Handbook — the Method the Inventor of the VSM Recommended for Using It

> **A note on standing.** Every other method in this series is someone's interpretation, extension, or application of the Viable System Model. This one is the model's author applying his own creation. Stafford Beer invented the VSM — he set out its scientific laws in *Brain of the Firm* (1972) and *The Heart of Enterprise* (1979). *Diagnosing the System for Organizations* (1985) is the third of that trilogy, and the only one written as a working manual: Beer's own step-by-step procedure for putting the model to use. When practitioners disagree about "what the VSM really requires," this is the primary source against which the disagreement is settled. It is the canonical method — not because it is the best for every situation, but because it is the originator's.

In Brief

Beer opens with a provocation aimed at the org chart. Our organizational structures, he writes, have been "frozen out of history" — built on the limitations our predecessors were forced to handle, then inherited unquestioned, so that even brand-new enterprises are "structured on the models with which we grew up." His alternative is not a better chart but a different object entirely: the scientific analysis of viable systems, meaning systems "capable of independent existence — in a word, surviving" (jacket text; Introduction).

Diagnosing the System is the handbook for doing that analysis yourself. It is explicitly "a Handbook or Manager's Guide" (jacket), pitched so that any manager can work through it "without requiring of him or her any prior knowledge at all" of cybernetics, using "less than twenty technical terms" (Introduction). The method reduces to eight teaching sections that build the model cumulatively, and — in the Appendix — four photocopiable **Charts** the reader actually fills in, ideally enlarged "several feet high" on a project-room wall.

This chapter describes that method, names what is unique about it (its provenance, its do-it-yourself design, and the parsimony of recursion), identifies its best application case, and states its limits — including its

considerable age. It then does what Beer himself declines to do: makes explicit the political assumptions of a method that addresses itself, in the same breath, "to managers and ministers" and treats autonomy as a matter of variety "rather than political theory."

Nine Managers and One Question

The method was forged in a classroom. In June and July 1982, Beer taught a post-graduate course at Concordia University in Montreal with Professor David Mitchell. His nine students had all already held management responsibility, and Beer asked them to tell him exactly where people got stuck with the VSM. The answer came back as a single recurring question: *"Yes, but how exactly do we proceed?"* (Introduction).

Brain of the Firm and *The Heart of Enterprise* had established the laws. What they had not done was tell a working manager what to do on Monday morning. So the class modelled seven viable systems together — a university, a manufacturing industry, the Quebec Health Service, a Third-World broadcasting corporation, a planetary-protection body, the Family, and language itself — and out of that came the procedure this book records. Beer dedicates the book to those nine students, and notes that their worked records stood "two metres high" (Introduction). The book is that course, reduced to something a reader can hold.

That origin matters for how the method reads. It is not a consultant's engagement model. It is a teaching sequence, addressed relentlessly in the second person — "the book is going to develop, with your intimate collaboration, a Model of the Viable System — the one that concerns *you* in particular" (Introduction) — and punctuated with boxed commands: NOW DO THIS. MAKE A LIST. Take a large sheet of paper.

The Method: Eight Sections, Four Charts

The teaching body walks the reader up the model one system at a time, each section closing with a short glossary of its new terms. The Appendix then supplies the worksheets. In practice the method has two layers: the sections teach you to see, the Charts make you commit what you see to paper.

Section One — the System-in-Focus and its recursions. The first task is to fix precisely which organization you are modelling, then name the viable system one level up that contains it and the viable systems one level down that it contains. Beer insists you always work with this *trio* — container, system, contents — because recursion is exact: every viable system is composed of viable systems and is itself contained in one, "upward to the terrestrial globe and downward to the cell." The central member of the trio is the **System-in-Focus**, and Beer warns from the worksheet itself that "it is

more difficult than people imagine to keep the System-in-Focus in focus" (Section One; Chart One).

Identifying System One. The single most error-prone step, and Beer says so: "most of the incorrect inferences derive from nominating activities that are not in themselves viable systems as if they were" (Section One). System One is the set of units that actually *produce* the enterprise — the parts that could in principle be hived off and survive. The discrimination test separates these producing units from mere *facilitators*: the sales department, the engineers, the computer department are worked examples of things that feel important but are not next-recursion viable systems.

Sections Two–Three — variety and the vertical axis. Here the reader learns to count complexity as *variety* — "the number of possible states of a system" — and to apply the Law of Requisite Variety: "only variety can absorb variety (Ashby's Law)." Variety is managed by **attenuators** and **amplifiers**, and every count is filtered by relevance: what matters is only whether a change of state bears on the system's purpose (Sections Two–Three). This is where the **command axis**, the **Resource Bargain** (the deal by which autonomy is granted in exchange for agreed accountability), and the **Accountability** channel are introduced — and where Beer plants the sentence that defines the whole model's stance: senior management is "in principle unable to entertain the variety generated by any one of its subsidiary viable systems. The beginnings of a theory of autonomy, of de-centralization, lie in this simple fact — rather than in political theory" (Section Three).

Section Four — System Two. The anti-oscillation function. Beer is careful that System Two "does not lie on the central command axis. Its function is not to command, but to damp oscillations" — the school timetable and production control are his canonical examples (Section Four).

Section Five — System Three, Three-Star, and the six channels. System Three is the "inside-and-now," day-to-day management of the whole cluster of operations. **Three-Star (3*)** is its sporadic, high-variety audit channel — "task-force" interventions into actual operations, of which "audit is a typical function." Beer enumerates the **six vertical channels** by which management absorbs operational variety, and states the **First Axiom of Management**: the sum of horizontal variety in all the operational elements equals the sum of vertical variety on the six channels of cohesion. The audit channel is what "closes the gap" (Section Five).

Sections Six–Eight — purpose, System Four, System Five. Section Six introduces autonomy proper — "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" — and the maxim "the purpose of a system is what it does." Section Seven adds System Four, "the outside-and-then," the organ of adaptation that must contain "a model of itself," balanced against System Three in the celebrated Three-Four homeostat (his example: the Rolls-Royce RB211 collapse). Section Eight closes the system

with System Five — identity, self-reference, logical closure — and the **algedonic signal**, the pain/pleasure alert that bypasses the ordinary channels: "the cry is 'wake up — danger!'" Each of the last three sections carries an Axiom, and the whole resolves into the **Law of Cohesion**, which Beer notes is simply the First Axiom restated in the language of the System-in-Focus (Sections Six–Eight).

The four Charts (Appendix). This is where the reader stops reading and starts diagnosing. Beer and Wiley grant a printed *certificate* permitting photocopying of the Charts alone, so the reader can "make up pads, or tablets, of each chart, so that rough analyses can be made, torn up, and improved upon."

- **Chart One** fixes the System-in-Focus and its recursions on a single sheet, to go up on the wall "preferably in colour, to keep all concerned alert to the set of recursions that has been chosen."
- **Chart Two**, spread over two pages, is the main diagnostic account of the viable systems at each recursion — corporation, then divisions, then companies — scaled by scissors-and-paste to the actual number of units. Beer's refrain: "it is the reality-out-there, and not the cybernetic technique, that generates the work."
- **Chart Three** handles the two channels readers most often get wrong — System Two and System Three-Star.
- **Chart Four** takes a *single homeostatic loop* and tests it against the Four Principles of Organization — checking that the attenuators and amplifiers around the loop actually deliver requisite variety. "When in doubt about the effectiveness of a homeostatic loop, analyse it with the aid of the facing Chart."

Diagnosis is meant to happen *as you go*, not at the end. Beer's instruction, riffing on his own epigraph: "make the diagnosis FOR ACTION as you go along; when you have recognized the remedy required, hold your nose, open your mouth, and swallow" (Introduction).

What Is Actually Unique Here

It is the source. Every interpretive question about the VSM — what counts as System One, whether a shared service is 3* or System Two, how recursion is bounded — has an answer here in the inventor's own hand. That is a standing this book has and no other VSM method can. It is also the least mediated: the fewest layers between the reader and the model.

It is genuinely do-it-yourself. The intended agent is the manager, working alone. "This book is meant to stand on its own, and to be fully accessible to the manager who will spare the time to work through it. There is, however, no reason why he should not lighten the task by taking a staff aide into his confidence" (Introduction). No consultant, no facilitation team, no prior cybernetics. Among serious VSM methods this is the most self-sufficient — and the cheapest to start.

It runs on the parsimony of recursion. Because every recursion has the identical structure, the analytic effort collapses. Beer claims very large systems — the Canadian fisheries, "the whole social economy of Chile (in Allende's time)" — were "modelled in less than two years, due to the parsimony of natural invariance." You learn one structure and apply it at every scale; "speedy diagnosis is a major contribution of the whole approach" (Appendix note).

It assesses structure against laws, not opinion. The Principles of Organization and the three Axioms give the reader variety balances to check rather than convictions to assert. In this it shares Pfiffner's instinct — measure, do not argue — but at the source, and more austerely. The laws are also the seed of the pathology catalogues that came later: what Hoverstadt (ch. 7) and Pérez Ríos (ch. 4) compile as named failure modes are recurring violations of the laws stated here — Beer's diagnosis *generates* from first principles what their catalogues *enumerate*, which is why his method can find a fault that has no name yet, and theirs can find a named fault fast. The lineage is documented, not inferred: Pérez Ríos cites Beer's own pathology paper — "The viable system model: its provenance, development, methodology and pathology" (1984; reprinted in Espejo & Harnden 1989) — as the analytical source of his catalogue.

Strengths, Best Case, Limits

Strengths. It is authoritative, self-contained, remarkably compact, and it forces action through its worksheet form. It teaches a whole modelling language in fewer than twenty terms. And it is honest about its own boundaries in a way few management books are.

Best application case. The method fits a manager or a very small group who want to understand and redraw the deep structure of an enterprise *themselves* — quickly, privately, and without hiring anyone — and who can tolerate abstraction and hand-drawn worksheets. It rewards the analytically inclined leader most, and it pays off disproportionately on large or bewildering systems, where the recursion parsimony turns an impossible amount of detail into a few repeated structures. It is also the right first stop for anyone who wants to check a claim about "what the VSM requires" against the primary source.

Limits — and Beer states most of them himself. The scope disclaimer is on the jacket: "It will not solve your financial problems, your personnel problems, or other managerial dilemmas. It is not concerned to spot major trends, nor in searching out slogans that might embody excellence." It deals only with "how to cope with complexity itself." Inside the book he is blunter still: "please do not expect this brief manual also to reflect on matters of psychology, of social anthropology, of industrial relations" — not because they don't matter, but "we cannot tackle everything at once" (Introduction). And there is a warning about reward: structural problems,

solved well, get *dissolved* rather than visibly fixed — "This means that you will not get any credit" (Foreword).

Two further limits are the reader's to weigh. The book is of 1985 and shows it: the Charts are hand-lettered, the imagined workflow is photocopier, scissors, paste and a wall; the examples are period pieces (the RB211, East–West détente, the Cold-War arms race); the default manager is "he." None of this touches the logic, but it dates the surface, and a modern reader must do the translation to software themselves. Second, the very self-sufficiency that is a strength is also a ceiling: a method designed for one diagnostician has no built-in way to create the shared understanding or the coalition that *implementing* a redesign requires — which is precisely the gap the later methods in this series were built to fill.

Part I — The Political Environment the Method Was Built For

Beer's explicit position is that the deepest political questions of organization are really questions of variety. Autonomy and decentralization "lie in this simple fact — rather than in political theory" (Section Three). Freedom is repeatedly re-described as residual variety: "the freedom remaining to the management on the horizontal axis to manage" (Section Six). Even litigation and demarcation disputes are to be pulled out of "an emotive and political matter" and restated "in cybernetic terms" (Section Five). This is the same technocratic move Pfiffner makes, and it is worth naming as a move rather than accepting as a fact: it is a claim that structural adequacy can be judged without asking whom the structure serves.

Yet the book's own framing is unmistakably political in ambition. It is addressed, on its first page, "to managers and ministers, their aides and their consultants," and its definition of "organization" runs "whether you are interested in a firm, an international conglomerate, a social service, a consortium of like-minded people, a government department, or a national economy" (jacket). Its epigraph is from a sitting head of government — Pierre Trudeau, 1969: "the many techniques of cybernetics... will transform our whole society. With this knowledge, we are wide awake, alert, capable of action; no longer are we blind, inert powers of fate" (epigraph). The claim is not modest. It is that the same laws govern the viable firm and the viable state.

And the method's provenance is, in fact, a state — and a particular one. The largest application Beer cites is the regulatory system he built for President Salvador Allende's Chile in 1971–73, "a decentralized real-time regulatory system for the whole social economy," described in full in the second edition of *Brain of the Firm* (Appendix; author biography). The originator's own flagship case is a democratic-socialist government's attempt to steer an economy without central planning — an emphasis, Beer stresses, on "the greatest autonomy possible" for each unit. The model

that presents itself as politically neutral was proven in one of the most politically charged experiments in the history of management.

Beer is not naïve about political conditions; he simply does not theorize them. His one aside on the difference is mordant. Recounting a year spent trying to apply the ideas in Mexico as he had in Chile, he remarks that "democracy in Mexico is not about casting votes, but counting them" (Introduction) — an acknowledgement, in passing, that the same method meets very different fates in different political soils, with no framework offered for why.

The best political environment for this method, then, is the one its provenance reveals rather than the one its rhetoric claims. It presupposes a commissioner who genuinely wants distributed autonomy — Allende's stated aim — rather than one who wants the appearance of it; a setting where the higher recursion's intervention rights are real but self-limited by design; and a leadership willing to accept that a viable design pushes decision-making *down* to where the variety is. Where those conditions hold, the method's silence on politics costs little. Where a commissioner wants cybernetic efficiency *without* the autonomy — the arms race and the surveillance state are among Beer's own variety examples — the model will still function, and, as with all the methods in this series, it will not by itself object to the use it is put to.

Part II — The Internal Power Constellation It Presupposes

Inside the organization, the method makes a single, sharp choice about power, and it is the opposite of the choice the other methods make: the diagnosis is performed by *one person*.

The manager holds the pen — alone, or nearly so. The book is written to be worked through by the manager, with at most "a staff aide taken into his confidence" (Introduction). There is no core team, no members' meeting, no facilitated workshop. This is the leanest possible power constellation: the person responsible for the System-in-Focus diagnoses it themselves. It presupposes that this person has both the authority to act on what they find and the intellectual appetite to build the model without help.

Power is acknowledged, then bracketed. Beer does not pretend organizations are free of it. He names "the existence in all viable systems of an equation of power," governed by "the will for survival" (Section Six). But he consistently refuses to let power define the systems: "System Three is not constructed as a box to house people with better suits and bigger cars" (Section Six), and System Five is emphatically not "the boss" — the metasystem is "over and beyond" a lower logical order, and "higher authority is not the issue" (Section Eight). Hierarchy, in this reading, is logical placement, not social dominance — the same reframing Pfiffner would later develop at length.

Participation is essentially absent, by design. This is the honest gap. Workers appear in the book almost only as a *source of oscillation* to be handled by System Two — "demarcation disputes between Trade Unions in every industry. Think hard, and beyond political prejudice, to reach the machinery of this" (Section Five). There is no participation mechanism, no consultation step, no account of how the people inside the System-in-Focus come to understand or accept its redesign. "Participation" occurs once in the whole book, describing the Concordia students. This is not an oversight so much as a boundary Beer draws deliberately — he excludes industrial relations and man-management explicitly — but the consequence is real: the method produces a diagnosis, not a coalition.

Where the internal preconditions fail. They fail wherever a redesign must be *agreed* rather than merely *seen*. A single sharp diagnosis is of little use in a partnership, a co-operative, a works-councilled firm, or any setting where legitimacy to change the structure is distributed rather than held. There, the diagnosis remains valuable — often the fastest route to clarity available — but it is a starting document for a process the book does not describe. This is exactly the seam along which the three methods in this series divide, and the division is clarifying:

- **Beer (1985): one diagnostician.** The manager works the charts alone. Fastest to insight; weakest on adoption.
- **Pfiffner (2022): a confidential core team of three to five,** designing in private for three months before a wider group of key people refines a design whose cornerstones are already fixed. A coalition, but a reserved decision.
- **Espinosa & Walker (2017): the whole community co-constructs** the model over months of workshops. Slowest; strongest on adoption and durability.

They are not rival readings of the model. They are three different answers to a question Beer's book does not ask — *who holds the pen?* — each right for a different distribution of power. Beer wrote the source. The others wrote the ways of using it when the manager is not, in fact, the only one who has to agree.

What to Borrow

Start from the source once. Whatever method you ultimately use, work Beer's four Charts once, by hand, on a real organization you know. It is the fastest way to internalize the distinctions every other method assumes.

The trio discipline. Never model a system without simultaneously naming the recursion above it and the recursion below it. Beer's insistence on the container-system-contents trio prevents the single most common error — analyzing a unit as if it had no context.

The System One discrimination test. Before anything else, separate the units that *produce* the enterprise (and could survive hived off) from the facilitators that merely support them. Beer's warning holds for every VSM exercise: most wrong conclusions start by mistaking a support function for a viable system.

Diagnosis for action. Do not save findings for a final report. Beer's "swallow the medicine as you go" is sound practice: each structural fault you identify should carry its remedy the moment you see it, because a diagnosis banked for later is a diagnosis quietly abandoned.

Read it to settle arguments. When practitioners disagree about what the VSM requires, this is the book that adjudicates — not because Beer is beyond improvement, but because it is his model, and this is the method he built for it.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1972). *Brain of the Firm*. London: Allen Lane. 2nd ed., Chichester: Wiley, 1981.
- Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley.
- Beer, S. (1984). "The viable system model: its provenance, development, methodology and pathology." *Journal of the Operational Research Society*, 35(1), 7–25. Reprinted in Espejo, R. & Harnden, R. (Eds.) (1989), *The Viable System Model: Interpretations and Applications of Stafford Beer's VSM*. Chichester: Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.

- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.

Chapter 2

The Third Dimension

Pfiffner's Seven-Step Method — and the Quiet Politics of a Well-Designed Hierarchy

In Brief

Organizations are designed in two dimensions. The first is **anatomy** — the org chart, essentially unchanged since Daniel McCallum drew one for an American railroad in 1855. The second is **physiology** — process design, which has been optimized continuously since business process reengineering in the 1990s. Martin Pfiffner's argument is that both have been exhausted, that a century of refining them has not produced satisfied managers, and that what is missing is a third dimension: **neurology** — what the organization does to manage its internal and external complexity. Control and communication. The Viable System Model, in his reading, is the only model that addresses it comprehensively (Preface, pp. v–vi).

What distinguishes his book from most VSM literature is that it is a manual. Seven steps, in order, each building on the last, with tools, tests and worked examples. And it makes one methodological move that the rest of the field mostly does not: where other approaches settle questions like *how much autonomy?* or *how flat can we go?* by argument, conviction, or fashion, Pfiffner settles them by **measuring variety**. "Wars of ideas, fads, and paradigms are unnecessary" (p. 104).

That is the strength, and it is also where the politics hides. This chapter describes the method, names its unique contribution and its best application case, and then does what the book itself does not: makes explicit the political assumptions the method carries — first at the level of society, where Pfiffner's implicit model is Swiss federalism, and second inside the firm, where his implementation design begins with three to five trusted people and no communication at all.

One Hundred and Seventy-Two Municipalities

About twenty years ago — Pfiffner leaves the date approximate — the Canton of Zurich held a constitutional debate. It has 172 municipalities reporting, in effect, to one cantonal level. That looked like far too many. The proposal

on the table was the obvious one: insert an intermediate control level between the municipalities and the canton (p. 155).

This is the kind of question organizations answer with rules of thumb — a manager can supervise seven people, or nine, or twelve — and then defend with anecdote. Pfiffner answers it by working the variety balance ("variety": the number of possible states) — a structured estimation on a small scale, it should be said, not a computation.

On the horizontal axis, three factors. There are 172 units, which is a lot. But they are barely differentiated: one municipality functions much like another, the same tasks and the same problems, with only the three larger cities materially different. And they largely solve their own problems, because Swiss direct democracy, subsidiarity and municipal fiscal sovereignty combine into a stabilizing negative feedback loop — the citizens who approve the indoor swimming pool are the citizens whose taxes rise (pp. 148, 155).

Then the canton's six vertical channels. Environmental overlaps: small, because regional boundaries are clear — though not negligible, since residents and businesses do not stay inside one municipality. System 3* (the sporadic in-depth audit channel): very strong, thanks to a broad media landscape through which the canton informs itself. Operational dependencies: strong — utilities, waste disposal, courts, hospitals, schools are shared at district level. Resource bargain and accountability: functioning, via fiscal equalization plus a legal obligation to coordinate plans. Corporate intervention: effective, because the canton can enforce its intervention rights. System 2: strong, supported by the districts, by law and ordinance, and by a common language, culture and religion (pp. 155–156).

The verdict follows from working the balance rather than from anyone's convictions about local self-government: the vertical axis is strong enough that the apparently unmanageable horizontal variety is, in fact, manageable. The extra level was not needed.

Whether or not one accepts the specific reading, the demonstration is the point. A structural question that had become constitutional politics was reframed as a measurement problem, and the measurement produced an answer that could be checked.

Three Dimensions, and Why the Third Is Empty

Pfiffner's diagnosis of the field opens the book (Ch. 2, pp. 25–48) as ten cardinal mistakes of organizing. Several are familiar: beginning with organizational charts ("Organizing does not start with organizational charts but ends with them", p. 26); copying competitors, since structure follows strategy and copying a rival's structure copies their strategy; escaping into the matrix; organizing past customer benefit. Two are worth pausing on because they set up the method.

Overlooking nesting. Managers routinely fail to say which recursion level they are speaking from — assessing a situation as head of an operational unit while sitting in a group-level discussion, or vice versa. Pfiffner returns to this repeatedly: "What level of recursion am I on when I say this?" (p. 105). The control functions are identical at every level; the content is entirely different.

Neglecting neurology. The tenth and organizing mistake. A new anatomy, he writes, "comes to life only when properly innervated" (p. 244). Reorganizations that redraw boxes without redesigning control and communication reliably shift friction from one place to another — which is why, as he observes, the number of organizational problems tends to be constant before and after a reorganization (p. 7).

Behind all ten sits the problem he considers decisive: we have no binding standard. "Lacking a reliable yardstick, we are disoriented and susceptible to fads" (p. 8). His candidate yardstick is complexity — measured as variety, the number of distinguishable states — on the grounds that it is the one common denominator across a car plant, a hospital and a newspaper publisher, and simultaneously their greatest challenge (pp. 9, 138).

The Method: Seven Steps

The diagnostic process (Ch. 7, pp. 115–120) begins "outside-in" and "bottom-up," at the customer and the operational unit — "the 'place of truth' and the pivot around which organizing turns" (p. 115). Pfiffner insists on *diagnosis* rather than *analysis*, since analysis breaks a system into elements and performance arises from their interaction. Diagnosis and design happen at once: "during diagnosis, we are inevitably already shaping the organization" (p. 115).

Before starting, four things must be known: the mission and any organizational options excluded a priori for corporate-policy reasons; the strategy, and whether it says anything about market, innovation, productivity and employer attractiveness; the culture and the key people; and the organizational history — what has been changed, when, and how successfully (pp. 116–117).

Then the unit to be designed is named the **System in Focus** (SIF) and labelled recursion level R0. Diagnosis always considers at least R+1 and R–1.

I — Basic structure (anatomy). Map operational units across all relevant recursion levels, ignoring support units for now. Check whether each unit can operate largely autonomously and whether it focuses on what customers pay for. Generate alternative segmentations — if the firm is currently split by product, test region, customer group, technology — and then select one, establishing the primary control dimension (pp. 117–118, Ch. 8).

II — Manageable or not? Test the span of control at every level by measuring variety rather than by rule of thumb, and test the maximum possible flattening: can a level be removed without overloading management? Output: the **IDEAL** structure across all levels (p. 118, Ch. 10).

III — Mission-critical tasks. Return to the SIF. For each operational unit, identify which tasks are mission-critical — the 20% on which 80% of success depends — including those arising from environmental overlaps and operational interdependencies (p. 118, Ch. 11).

IV — Central or decentral. Across levels, decide which task belongs where. Note the stated order of priority: first establish which competencies the operational units *must keep* to maintain relative autonomy and satisfy their customers; "only secondarily do we examine which tasks can be centralized in order to achieve synergies" (p. 119).

V — Control functions. Distribute those tasks across Systems 2, 3, 3*, 4 and 5. Typically 20–40 tasks are mission-critical at a given level. Then assign responsibility for each control function — and here Pfiffner makes a sharp practical observation: someone is normally named for System 3 (usually the CEO) and often for System 5 (board chair or president), but "for control systems 2, 3*, and 4, on the other hand, no responsible persons are designated as a rule — which explains why these systems are often neglected in practice and cause control problems" (p. 119).

VI — Communication channels. Test the channels between the functions for capacity, intelligibility, timeliness and security; connect the levels; and check whether the algedonic signal (Beer's pain/pleasure alarm bypass) actually works (pp. 119–120, Ch. 14).

VII — Making it understandable. Only now do positions and people enter. A control function diagram represents the whole control system of the SIF on a single page, from which org charts and job descriptions are derived (p. 120, Ch. 15).

On scale, the book is deliberately elastic and correspondingly vague: diagnosis "may take up 5 min in our own mind, be developed on a napkin over dinner for two, or require some executive retreats over several months" (p. 116). Team size is not vague: two to five people, "compromising between ensuring the necessary productivity and enough variety for good discussions" (p. 116).

What Is Actually Unique Here

Measurement in place of ideology. This is the book's signature and its real contribution. Autonomy versus hierarchy, centralization versus decentralization, flat versus layered — Pfiffner treats each as a computable balance for a specific situation rather than a position to hold. "There is no clear 'the more the better' and no ideological solution but only a suitable balance for the current situation" (p. 104). Whether one finds his variety arithmetic rigorous or merely structured, it changes the character of the argument: the claim becomes checkable.

Hierarchy re-derived from information rather than status. His treatment (Ch. 6, pp. 101–112) is the most useful defence of hierarchy in the VSM literature, precisely because it does not defend the thing usually meant by the word. "In a viable system, hierarchy results from the relevance of information (rather than from status and power)" (p. 102). The pilot outranks the air traffic controller in status and salary and obeys the tower on approach, because at that moment the tower holds the relevant information. "The command center always shifts to the place having the most information" (p. 102). Senior management is then "a meta-systemic service provider," and the top manager "not the most important person but the enterprise's highest servant" (p. 103). Autonomy is restricted for exactly two legitimate reasons: to protect the identity of the whole (System 5) and to realize synergies that make everyone stronger (System 3) — with Volkswagen's platform strategy and the City of Zurich's incompatible departmental IT as the worked examples (p. 104).

The IDEAL-before-TARGET sequence. Borrowed from Walter Krieg of the St. Gallen school. The path from the CURRENT to the implemented TARGET organization runs via an IDEAL organization that is correct *ad rem* — right for the matter, designed only from purpose and customer, with no consideration of who is available to fill the roles. Only then is it tested *ad personam*, and compromises are made. The point of constructing an ideal one knows will be compromised is to be able to distinguish good compromises from bad ones: "From all possible compromises, we need to choose that one that most closely approaches the ideal" (p. 238).

A completeness test. Because the five systems are necessary and sufficient, the model tells you when you are *finished* organizing — not merely when you are tired. In a field with no standard, that is not a small thing.

Strengths, Best Case, Limits

It is executable. Numbered steps, an order that must be broadly respected because each builds on the last, named deliverables, a function diagram that fits on one page, and a chapter of seven named pathologies with symptoms and remedies (Dissociation, Schizophrenia, Bottle-Neck, Uncontrolled Growth, Dominance, Multiple Personality, Mindlessness — Ch. 17, pp. 247–270) for readers who want a fast read on a familiar situation.

It fixes what most reorganizations miss. The observation that Systems 2, 3* and 4 have no owner is a finding of real and lasting use. Most firms can name who owns operations and who owns policy; almost none can name who owns coordination, real-life audit, or outside-and-then (System 4, the outside-and-future function).

It de-escalates. Turning "should we be flatter?" into a variety calculation removes a class of unwinnable meeting from the calendar.

Best application case. The method is at its strongest in a mid-sized to large enterprise, under a single accountable leadership, facing a genuine structural question — a segmentation that no longer matches the market,

an acquisition to integrate, a span of control that has grown past comprehension — where the leadership can decide and implement, and where a nine-month reorganization is an acceptable instrument. Its natural client is the executive who suspects the org chart is the wrong tool and has no better one. That is precisely who Pfiffner writes for: "This book is for those managers who want to put Beer's model into practice" (p. vi).

Limits, stated plainly. The complexity "measurement" is structured estimation, not quantification — you are ranking factors on a small scale and weighing them, not computing a number, and the book does not pretend otherwise, though a fast reader may miss the distinction. The claimed elasticity of scope (five minutes to several months) makes effort genuinely hard to plan. The pedagogical device of recurring fictional dialogues (Rachel, Marc, Sandra, the pilot David) will help some readers and irritate others. And the book is a practitioner's manual, not a research report: the evidence is two decades of consulting experience, presented as illustration rather than as tested claim.

Part I — The Political Environment the Method Was Built For

Pfiffner's explicit position is that the method has no politics. Structural questions "independent of ideological ideas or personal interests" (p. 104); successful examples of both hierarchical and autonomous control exist; "Whether something works does not depend on whether it is hierarchically or autonomously controlled but on whether the control system per se works" (p. 101).

That stance is itself a political position — a technocratic one, in the precise rather than pejorative sense: the claim that structural adequacy can be assessed independently of who benefits. It is a defensible claim and it is not a neutral one, because it removes questions of interest from the frame by construction. When Pfiffner writes that hierarchy discussions are "as a rule no more than an ideologically motivated dispute" (p. 103), the possibility that a dispute about hierarchy might be a real conflict over real interests, correctly perceived by both sides, is not so much rejected as set aside.

Underneath the stated agnosticism, an implicit political model is visible throughout, and it is Swiss. The examples that carry the argument are subsidiarity, fiscal sovereignty, direct democracy, federalism, and free media. His account of why Swiss public finances are healthy is that "not the typical Swiss virtues keep the country financially healthy but combining autonomy with a stabilizing feedback loop" — those who vote for the swimming pool pay for it (p. 148). His general formulation of this is worth quoting for how much politics it compresses into one line: "Change becomes easier, the better we not only turn those affected into participants but also turn participants into those affected" (p. 148).

The model extends to the state without strain. Municipalities are the lowest recursion level; districts, cantons, countries, Europe above them (p. 94). The executive is System 3, parliament System 4, the constitution and constitutional court System 5, with citizens deciding at the polls and — in direct democracy — on substantive questions such as fighter jets, tax law and whether cows should have horns (pp. 85–86). Free media perform both a System 3* function (unfiltered ground truth reaching the top) and a System 4 function; "This is probably the systemic reason why media freedom is so highly esteemed in democracies" (p. 83).

One passage deserves attention because it is uncomfortable and Pfiffner does not soften it. Discussing System 4 and the vertical axis, he notes that European states change "much more sluggishly than, for example, in China at the beginning of the new millennium, where a stronger vertical axis of control is evident," and that many countries fail to make adjustments whose necessity has long been visible — demographic ageing being his example (p. 83). He is describing an adaptation-speed property of a control structure, not endorsing a regime. But the book contains no counterweight — no discussion of what a strong vertical axis costs when it is wrong, and no acknowledgement that variety can be attenuated by coercion — by locking people up — which is why requisite variety alone is no ethics (Espinosa and Walker at least touch on this; see ch. 3). A reader inclined to read cybernetics as a technology of control will find nothing here that argues them out of it.

The most politically revealing moment is an anecdote. Beer explained the model to Salvador Allende, and as he reached System 5 and was about to tell the president that *he* was System 5, Allende exclaimed: "Ah, finalmente — el pueblo!" Pfiffner's gloss: "the lowest authority in the organization, the citizen, toward whom the entire state is directed, is also the highest authority. This need not be the case in every state but it has interesting, self-regulating effects" (pp. 87–88). The claim that System 5 is the people is treated as one design option among others, notable for its regulatory properties. The same anecdote in Espinosa and Walker would be a normative commitment.

The best political environment for this approach, then, is one Pfiffner describes rather than argues for: a federal, subsidiarity-based polity with genuine local fiscal autonomy, working feedback between decision and consequence, free media supplying the audit channel, and intervention rights at the higher level that are real but bounded. Where those conditions hold, the method's agnosticism costs nothing, because the political question has already been settled well. Where they do not — where the higher level can intervene without limit, where those who decide do not bear consequences, where there is no independent audit channel — the method will still produce a structurally sound design, and will have nothing to say about the fact that the design serves whoever commissioned it.

Part II — The Internal Power Constellation It Presupposes

Inside the firm, the method's assumptions about power are unusually explicit, and reading them carefully matters more than reading the seven steps.

It is commissioned from the top and it stays there. The reader addressed is an executive responsible for a recursion level — "we start on that level of recursion that we can actually design because we are responsible for it" (p. 118). Where the book names owners of control functions, they are the CEO (System 3) and the board chair or president (System 5) (p. 119).

The first three months are confidential. This is the sharpest structural commitment in the book and it should not be skipped. Phase A: three months, a core team of three to five trusted people, communication — *none* (Table 16.1, p. 240). The core team develops the IDEAL organization and tests it *ad personam*. If the conclusion is to stop, "nothing should be communicated up to this point, because communication would only unnecessarily unsettle the enterprise" (p. 240). Phase B: four days, the core team plus 20–40 key people, with a brief announcement to the company. Phase C: three months, core team again, updates only to those key people. Phase D: six months, implementation, detailed information to everyone. Total: nine months, "in employees' perception, from first mention to completion" (p. 239).

The cornerstones are not up for discussion. When the 20–40 key people arrive in Phase B, the segmentation and the number of recursion levels — the output of Steps I and II — are declared binding. "we do not want to encourage discussion of the cornerstones of the new organization… We define these as binding" (p. 241). Within that boundary, dissent is actively wanted: "we are looking specifically for dissent, because sustainable solutions emerge from enacted dissent rather than from hasty consensus," and the key people "may question everything — provided they propose better ideas" (p. 241).

Participation is instrumental, and the book says so. The mechanism for Phase B is Beer's Syntegration, which Pfiffner describes with evident enthusiasm: the icosahedron gives a structure that "is nonhierarchical, because a globe has neither top nor bottom. It is also perfectly democratic, because every participant has the same possibilities to influence the outcome" (p. 277). And then, immediately: "Whether the democratic form is also used for democratic decisions is an open question. In most enterprises where I have applied this method, the board reserved its right to decide but wanted to know the integrated opinion of its key people before deciding" (p. 277). He reports the outcome distribution: roughly 70% of proposals accepted, 20% revised, 10% rejected — his own recollection of his consulting practice, offered as illustration rather than audited data — adding that if key people saw the world very differently from the board, "something would be fundamentally wrong with the enterprise."

This is a democratic *process* operating inside a reserved *decision*. Pfiffner is entirely open about it, which is more than most participation formats manage. It is also the exact point at which this method and the participatory VSM traditions part company: the same instrument, with the authority left in a different place.

People come after things, deliberately. Steps I and II are performed while forcing oneself "not to think about specific people at first" (p. 238). Only then: do we have enough managers, who will lose out, and with which key people must solutions be found in personal discussion. Where no solution is found, the IDEAL cannot be implemented "or at least not yet" (p. 239). This is a defensible sequence and an uncomfortable one, and it presupposes a leadership willing to conclude that the right structure costs specific named people their positions.

Implementation runs on authority, exercised well. "Powers Reserved" — what is not forbidden is permitted, with reservations defined per level and per unit, everything else decided subsidiarily (p. 244). Adapt the incentive system immediately to the new dominant control dimension, on the principle of "control what controls" (pp. 243–244). Communicate directly, bypassing hierarchy, never in a cascade. Build an algedonic channel and a System 3* into the implementation project itself so the steering committee learns what is not working (pp. 242–243). Throw the lever in one move: "If you want to introduce something new, you must do away with the old for good" (p. 242).

What is absent. Works councils appear twice in the book, in passing, as possible System 2 or System 5 participants. Unions appear once, in a list of management tasks. Co-determination — which is legally binding for reorganizations of this scale in Germany and Austria, and consequential in Switzerland — is not treated. Nor is the question of whether nine months of a confidential core team is compatible with consultation obligations. This is a real gap rather than a stylistic one, and a reader in a co-determined jurisdiction will have to bridge it themselves.

Where the internal preconditions fail. They fail where no one person or body is accountable for the SIF; where the organization is a genuine partnership, co-operative or federation whose members will not accept binding cornerstones set by three to five people; where a nine-month reorganization is unaffordable; and where legal consultation requirements make Phase A's confidentiality untenable. In those settings the diagnostic content remains fully usable — the recursion mapping, the segmentation test, the ownerless-control-function audit, the channel checks — and only the implementation design has to be replaced.

What to Borrow

The ownership audit. Ask who owns System 2, System 3* and System 4 by name. If the answer is nobody, you have found the reason coordination is informal, the top has no unfiltered ground truth, and the future is discussed only at the annual offsite.

The recursion-level discipline. Require people in structural discussions to state which level they are speaking from. It is cheap and it removes a surprising share of circular argument.

IDEAL before TARGET. Design the structure the business needs before considering who is available. Then compromise consciously, and keep the ideal written down so that you know which direction later increments should move in.

"Powers reserved." Define reservations explicitly per unit and let everything else be decided subsidiarily. Pfiffner's argument for the asymmetry is practical: an overstep is visible quickly, whereas someone who does not use their scope "may not [be noticed] until after much valuable time has been lost" (p. 244).

The nine-month ceiling. Whether or not you accept the confidential opening, the constraint that a reorganization must not exceed nine months in employees' perception is a good one, and violating it is the standard way reorganizations become permanent.

What should be borrowed knowingly rather than by default is Phase A. Three months of confidential design by three to five people is coherent, defensible, and a specific choice about who is a participant and who is a recipient. Espinosa and Walker's method makes the opposite choice and both work in their respective settings. Choosing without noticing that a choice is being made is the failure mode — and it is the one Pfiffner's own framing, with its insistence that structure is a technical rather than a political matter, makes easiest to fall into.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.

- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.

Chapter 3

The Map They Draw Themselves

Espinosa and Walker's Self-Transformation Method — and the Politics It Quietly Requires

In Brief

Most structural interventions follow the same shape: an outside expert studies the organization, produces a diagnosis, and presents a target operating model. The organization then attempts to implement someone else's understanding of itself. Adoption is usually the point at which the effort dies.

Angela Espinosa and Jon Walker built a method that inverts this. In their *self-transformation methodology*, the consultant never delivers a diagnosis. The consultant teaches the Viable System Model as a **meta-language** — a shared vocabulary for talking about complexity — and then facilitates the organization's members in producing the diagnosis themselves, in workshops, over months. The model is not applied to the organization; it is co-constructed with it (pp. 94–98, Appendix A).

This chapter describes the six stages of that method, names its three real strengths, identifies the situation where it clearly outperforms alternatives, and isolates its genuinely unique contribution. It then addresses a question the method's advocates tend to skip: the approach was designed in and for a particular political environment — co-operatives, eco-communities, participatory state programmes, commons — and it carries assumptions about power that do not travel silently into a shareholder-owned firm. Those assumptions are worth making explicit, because they determine whether the method will work for you or merely produce a very well-attended failure.

Twenty-Two Working Groups and No Meta-System

In summer 2007, an Irish eco-community invited Espinosa and Walker to help. The project — a hundred-plus houses on a green-field site, super-insulated, reed-bed wastewater treatment, wood-and-solar district heating — was in trouble. The housing market had crashed, cutting the capital members could raise. A civil engineering survey had been inadequate and site costs had risen. Infrastructure was late (pp. 338–339).

The organization was, by any conventional measure, admirably democratic. A hundred-plus members met monthly and decided policy by consensus. Twenty-two self-organized working groups had formed around members' interests and skills. A board was elected. There was even a standing *process group* whose job was to reflect on how the organization worked.

And it was barely functioning. Nothing finished on time. The board kept stepping into operational work because the work was important and wasn't getting done — which the working groups experienced as interference (pp. 341–342).

The first workshop did not begin with a diagram. It began with cartoons. Members were asked to draw the problems they were experiencing (p. 340; the technique is Checkland's rich pictures, and the stated reason for the humour is that it moves a stuck group out of defensiveness). Only then were the VSM distinctions introduced, and the twenty-two groups mapped onto them.

The mapping made two things visible in an afternoon. There were far too many System 1s to be held together coherently. And there was no focused meta-system at all — Systems 2, 3 and 4 were unclear, and the board was simultaneously operating, managing and setting policy.

The facilitators then did something that looks small and is actually the whole method. Rather than reorganize the twenty-two existing groups, they asked the members a question and let them answer it: *which activities, between them, actually build the village?* The members produced seven: green infrastructure; site building and maintenance; site sales; enabling members to build houses; building community houses; education and dissemination; creating a sustainable community (p. 342).

Nobody had handed them that list. By 2010, essentially all the recommended changes had been implemented: seven primary activities, each with a coordinator; a coordination meeting with a standing System 4 slot; a named System 4 role for the whole community; the board finally confined to policy (pp. 343–345). Financial control was recovered during 2009.

The largest obstacle along the way is worth recording, because you will meet it. The members feared the VSM would impose "a straight-jacket" on their structures. It took months of internal debate before the changes landed, and the objection dissolved only when the facilitators explained that the model supplies design constraints, not a design (p. 343).

The Method, in Six Stages

The methodology is documented in Chapter 2.5 and set out operationally in Appendix A (pp. 473–484). Espinosa and Walker are explicit that the stages are "systemically connected, rather than systematically connected" — you run a full first cycle, then revisit stages as needed rather than marching through them (p. 95).

1. Understanding organizational identity. Desk research on the organization, then a workshop with representatives from everywhere. Cartoons first. Then a root definition using Checkland's CATWOE or Espejo's TASCOI: what transformation does this organization actually perform, and who are the actors, customers and owners of it? (pp. 473–474)

2. Recursive analysis. Separate primary activities — those that directly produce the organization's purpose, have a boundary, resources, local management, and could in principle exist independently — from support activities. Then unfold: draw the organization as a circle, its primary activities as circles inside it, and zoom into each until no further primary activities exist. Finally, agree the *system in focus* (pp. 474–477).

3. VSM analyses. For each parent–child pair in that recursion tree, map S1, connect S2–S5, draw the environment, classify roles. Three sub-analyses: the S1–environment variety engineering (where are the recurrent breakdowns?), the monitoring-control mechanism (S1–S3 — corporate regulation, resource negotiation, accountability, conflict resolution, monitoring), and the adaptation mechanism (S3/S4/S5 — which committees, councils and forums actually make policy, and is the S3–S4 balance real?) (pp. 477–481).

4. Aligning strategy and structure. Agree critical success areas and factors, then check that the system in focus actually contains them. Optionally, strategic information management (pp. 481–482).

5. Agreeing and implementing changes. Negotiate the findings with those who can allocate resources; establish what is feasible and in what order (p. 482).

6. Monitoring performance. Rebuild the performance management system on VSM lines: indicators mapped to primary activities and recursion levels, seven-plus-or-minus-two per unit, actuality/capability/potentiality rather than historical reporting, algedonic alerts rather than monthly packs (pp. 483–484).

What Is Actually Unique Here

Three things distinguish this from other VSM methodologies — Espejo's VIPLAN, Schwaninger's Systemic Management Control, Pérez Ríos, Hoverstadt — all of which share the same underlying model.

First: the analyst refuses to diagnose. The role is explicitly "a facilitator, a coach, who will provide some learning tools" (p. 95). The stated purpose of the structural work is *not* to produce an ideal model from the theory's point of view, but "to facilitate the learning of those participating in the exercise" (pp. 96–97). This is a genuine methodological commitment, not a soft framing: the target of the intervention is the organization's capacity to model itself, and the structural improvements are a by-product. Espinosa and Walker report that the VSM distinctions become part of the vocabulary and get reapplied to new problems after the consultants leave — which, in the eco-village, they demonstrably did (pp. 96, 345).

Second: viability is extended to the ecological niche, and sustainability is redefined as long-term viability. This is the book's central theoretical move. Sustainability stops being a separate programme with its own department and its own reporting, and becomes a property of the S1–environment loop at every level of recursion: what the unit draws from and returns to its niche, measured by essential variables the participants choose themselves (pp. 118–120, 331–332). The practical consequence is that "working out what matters" — carbon, resource depletion, biodiversity, human well-being, alongside financial measures — becomes a variety-engineering question rather than a reporting question.

Third: the method is scale-agnostic by construction. The same six stages are documented at the scale of a worker co-operative (Suma), a Mexican firm, a university, a national supply network, an eco-region river basin, a national anti-poverty programme, and — in Chapter 7 — an individual household. This is not opportunistic breadth; recursion is what makes it coherent.

Strengths — and Where They Come From

It survives the consultant's departure. Because participants build the model, they retain the language. This is the method's strongest empirical claim and the eco-village is its best evidence.

It reaches decisions that no diagnosis can force. The seven primary activities were an act of self-definition, not a recommendation. Recommendations get accepted or rejected; self-definitions get defended.

It handles genuinely fuzzy organizations. Volunteer communities, networks, co-operatives and multi-agency programmes have no reliable org chart to analyse. The recursion exercise builds the structure from what the organization *does*, which is the only thing such organizations reliably have.

Best application case. The method is at its clear best where four conditions hold together: the organization is complex enough that no one person holds a working model of it; ownership of the structure sits with the people inside it rather than with an external principal; the timeframe permits several months of workshop cycles; and the organization already has *some* legitimate forum where a whole-system decision can be made and stick. The eco-village had all four. Where the fourth is missing, the method produces excellent shared understanding and no structural change.

Honest limits. Espinosa and Walker's own account of the difficulties (pp. 385–386) is short and names three: an autonomous unit's sustainability is conditioned by the system it is embedded in, and "it is not something each embedded organization may achieve independently"; where there is no local leadership or resource, sustainability questions simply queue behind more urgent ones in a bureaucracy; and the whole thing presupposes "an ethos of co-operation rather than competition." They do not discuss cost, facilitator dependence, participation fatigue, or the possibility that

participation gets captured. Their own Colombian case notes that "the quality of participation was not always good" and that social control schemes "worked well in some municipalities but not in others" (p. 377).

Part I — The Political Environment the Method Was Built For

The book is not politically neutral and does not pretend to be. Reading it as a value-free toolkit will produce misunderstanding in both directions.

Chapter 6 opens with Naomi Klein and proceeds to a direct critique of the mainstream development paradigm: the Brundtland-to-Earth-Summit consensus as objectives without implementation (p. 321); the World Bank's "environmental economics" chain of free trade → growth → less poverty → less degradation as something that "has been shown not to work in practice" (p. 322); the post-1990 neoliberal model, named as such and traced to Friedman and the Chicago School (p. 323). Development expertise is described as top-down by construction — "the technical experts from multinational agencies assumed they knew better than the local people" — and conventional participation as a mobilization device: "The community was consulted but had no involvement in the creation of the plans or in their implementation" (p. 322).

From that critique the authors derive a positive requirement rather than a slogan. Participation without decision rights is worse than no participation: decision-making spaces "must have real decision-making powers — the 'consult and ignore' approach is worse than remote decision-making" (p. 335). And the political precondition for the method is stated plainly: each community "needs to be allowed to evolve independently of repressive authoritarian governments that would stifle self-organization and emergence" (p. 331).

The environments where the method demonstrably works share a family resemblance. Worker co-operatives — Suma, Mondragón, the Emilia Romagna and Cleveland Evergreen co-ops. Eco-communities. Transition Towns and the Transition Network. Long-enduring commons, analysed at length through Elinor Ostrom (pp. 411–424). And, conditionally, state programmes: Beer's Cybersyn under Allende, and Colombia's Social Solidarity Network, which the authors attribute to "a very positive coincidence of political will," technical expertise and local acceptance (p. 367).

What unites these is a specific structural feature the authors identify in the commons work: in a successful CPR, System 5 is invariably a meeting of *all* users, so System 3, System 4 and System 5 are articulated by the same assembly. Compliance is then high despite weak sanctions, because "participation engenders an identification with the meta-system" (pp. 421–422). That is a governance form, not a management technique.

The book is equally clear about hostile environments, though it frames them as barriers rather than exclusions. Command and control: "the vast

majority of businesses still works on the command and control paradigm and, as we have argued, are unlikely to be able to make [environmental management] policies respond quickly and effectively" (pp. 246–247). Shareholder primacy: "unless this changes it seems unlikely that much else will change" (p. 246). And the sharpest observation in the book about why the model spreads slowly — recursion implies that the complexity a manager handles is broadly similar at every level, which makes 300:1 pay ratios hard to justify, and this, the authors suggest, "is perhaps a reason the VSM has not been more widely adopted" (pp. 189–190).

Two conjunctures deserve attention because they show the method's political dependency in both directions. Cybersyn worked — the 1973 Gremio strike was broken by real-time routing of the remaining vehicles, and a senior minister said the government would have collapsed without the cybernetic tools (pp. 147–148) — and then ended with the coup. And in Colombia, the eco-region work was "abruptly stopped" by a change of government, with the agencies reverting to their previous practice; the authors conclude that the design was possible "but that politics can intervene and re-establish the old (non-sustainable) methods of working" (pp. 267, 428).

The best political environment for this approach is therefore not "progressive government." It is a setting where the group doing the work also holds the authority to change its own rules, and where that authority is not revocable by a superior body on a shorter cycle than the transformation itself. Ostrom's Principle 3, in the authors' gloss: "the rules are changed by the people who are affected by them" (p. 422).

Part II — The Internal Power Constellation It Presupposes

The societal argument matters less to most readers of this book than a narrower question: what must be true inside *your* organization for this method to work?

A sponsor who can be bound by the outcome. The method produces conclusions the sponsor did not commission — in the eco-village, that the board should stop doing operational work. If the sponsor can veto findings without cost, participants learn this within two cycles and the workshops degrade into consultation theatre. This is the "consult and ignore" failure moved indoors, and it is the single most common way the method fails in corporate settings.

A legitimate whole-system forum. Somewhere a decision must be reachable that binds all the primary activities. In a co-operative this is the members' meeting. In a company it is usually an executive team — which is fine, provided that team is in the room for the workshops rather than receiving a readout.

Willingness to name primary activities against the org chart. The recursion exercise routinely reveals that a unit everyone treats as core is a

support activity, and that a real primary activity has no home. Both findings redistribute status and budget. If that redistribution is off the table, stage two cannot be completed honestly, and everything downstream inherits the compromise.

Tolerance for the resource bargain. Devolving control means intervening only when a unit is out of control at its own level and has not corrected it within an agreed period. Management by exception, in other words — which requires managers to give up routine visibility. Espinosa and Walker's own Suma case records that VSM performance indicators were initially seen as a "Big Brother" control system, and that the objection dissolved only because each work group selected its own indicators (pp. 201–202). Who chooses the measures determines whether the same instrument is surveillance or self-regulation.

Where the internal preconditions fail. They fail where structure is the CEO's prerogative and not a subject for collective inquiry; where a matrix makes the boundaries of primary activities genuinely undecidable; where the reorganization horizon is a quarter and the method needs three; and where the honest answer to "what does this organization do?" is politically unsayable. In those settings the diagnostic parts of the VSM remain fully usable — the transmission gap analysis of Chapter 5, the resource bargain, the S3–S4 balance check. What is not available is *self*-transformation, because the self in question does not hold the pen.

What to Borrow If You Are Not an Eco-Village

Four elements transfer cleanly into ordinary corporate conditions, in rough order of ease.

The cartoon opening. Before any structural diagram, have the group draw the problems. It is not a warm-up; it surfaces the material the later analysis needs, and it does so before people have found their positions.

The primary-activity question. Ask a leadership team, without reference to the org chart: which activities, between them, actually produce what we sell? Expect the list to differ from the reporting structure, and expect the discussion of the difference to be the most useful hour of the quarter.

Indicator ownership. Let each unit select the small set of indicators by which it will be held accountable, subject to a constraint on how many and on what they must cover. This is the cheapest available intervention with a genuine effect on whether measurement produces self-regulation or defensive reporting.

The adaptation-mechanism inventory. List every committee, council, board and forum where sensing (S4) and current operations (S3) meet to set policy. Then ask, for each: is the balance real, or does one side arrive

with information the other cannot contest? Chapter 15's Adaptive Space is the same question asked from the other end.

What does not transfer is the assumption underneath all of it — that the people doing the work also hold the authority to redesign it. Where that assumption holds, this is the most durable VSM methodology in the literature. Where it does not, borrowing the tools is sensible and calling the result self-transformation is not.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5.

Chapter 4

The Method With a Machine

Pérez Ríos's Viable System Method — Diagnosis by Checklist, Matrix, and Software

In Brief

Beer wrote the model and a handbook for using it by hand. Espinosa and Walker turned it into a language a community speaks about itself. Pfiffner turned it into a seven-step design process. José Pérez Ríos did something none of them attempted: he turned the Viable System Model into an *instrument* — a fully specified procedure, exhaustive enough to be followed by an analyst who is not a cybernetician, and backed by purpose-built software that draws the model, navigates its recursion levels, and stores the entire study.

His diagnosis is that the VSM has spread too slowly for a specific, fixable reason: it is "difficult both to understand and to apply", and it has lacked "specialised software to act as a guide in its application" of the kind that carried Forrester's System Dynamics into wide use (Preface). His book is the remedy — the most step-by-step VSM manual in the literature, organized around three distinctive instruments: a four-stage procedure with a triple test applied to every element; the **Recursion Levels–Critical Factors Matrix**; a systematic **catalogue of organizational pathologies**; and **VSMod®**, the software he built at the University of Valladolid to make multi-criteria recursion tractable on a screen.

This chapter describes that method, names what is genuinely unique about it, identifies its best application case, and states its limits — including its dependence on a specific piece of software. It then makes explicit the political assumptions the method carries. Two are worth flagging at the outset, because the title invites misreading: for Pérez Ríos, "sustainable" means *viable* — long-term survival — not ecological, and the method is expert-led analyst's desk work, with participation confined to a separate, optional chapter.

A University and Its Territory

The worked example that runs through the book is not a company. It is a public university embedded in a region: the University of A Coruña, in

Galicia, and the town-planning question of how a university should be organized across the territorial scales it actually inhabits (from the study with Martínez, used throughout Chapters 2 and 4).

Pérez Ríos builds the recursion structure from the territory outward. Galicia is one scale; the A Coruña–Ferrol urban region another; the individual campuses another. Each scale has its own environment — its own institutions, legislation, stakeholders, and relevant issues — and therefore, in VSM terms, needs its own viable organization, not a branch office of the one above. He fills a matrix, level by level and column by column, and the exercise ends by generating seventeen concrete urban-planning actions (around printed p. 100).

The choice of case is not incidental, and it tells you what the method is *for*. Where Pfiffner reaches for Volkswagen and the Canton of Zurich, and Espinosa for an eco-village, Pérez Ríos reaches for a public university and its region — a system with multiple legitimate ways to cut it into levels (by territory, by academic discipline, by administrative function), no single obvious hierarchy, and a genuinely complex environment. That is precisely the situation his method is built to handle, and precisely the situation that defeats a pen-and-paper approach.

The Method: Four Stages, One Triple Test

Pérez Ríos states the procedure as four stages up front (Chapter 2, printed p. 65), applicable identically whether you are *diagnosing* an existing organization or *designing* a new one.

Stage 1 — Identity and purpose. Make explicit what the organization *is* and *is not*, and what it is for. This fixes the boundary between organization and environment. Note a deliberate divergence from Espejo's VIPLAN tradition: Pérez Ríos does *not* use the TASCOI mnemonic. His identity step is his own boundary-and-purpose formulation, and the full treatment of identity is deferred to System 5 later in the analysis (printed pp. 66–69). The step then characterizes the environment, split into **present** (clients, suppliers, legislation, competitors, institutions) and **future** (technological change, new markets, ecological restriction), across a checklist of environmental areas — economic, sociological, political, legislative, institutional, market, supplier, competitor, technological, ecological, educational, demographic — and defines the *sensors* by which the organization will detect each (printed pp. 67–69).

Stage 2 — Vertical breakdown. Choose one or more **recursion criteria** and unfold the environment and the organization into nested levels, each of which is itself a complete viable system. This is where his framing is most distinctive: an organization can legitimately be cut into levels several different ways at once — by territory *and* by product *and* by function — and the system-in-focus sits at the *intersection* of these criteria, in what he treats geometrically as a sphere rather than a single ladder. The image is worth a moment: picture each recursion criterion as an axis running through the

organization — territorial levels one way, product lines another, functions a third. A single criterion gives you a ladder of levels; several at once give you a space, and any system-in-focus is a point in it, reachable by more than one path of descent. Beyond two criteria this becomes impossible to draw on paper, which is his standing argument for the software (printed pp. 70–77).

The Recursion Levels–Critical Factors Matrix (printed pp. 78–79) is the core deliverable of this stage and the method's signature worksheet. Rows are recursion levels; columns are the critical factors to be established for each. He offers a ten-column "orientative decalogue": the level's description; its spatial scope; its relevant issues and purpose (congruent with the whole); the organization operating there; its stakeholders; the influential external institutions; the applicable legislation; the actions (chosen, in design mode; evaluated, in diagnosis mode); the means (what, who, how, when, where, cost); and the communication channels. The analyst fills it cell by cell. It is the artifact that forces the vertical analysis to be complete rather than impressionistic.

Stage 3 — Horizontal study. For the chosen system-in-focus, analyze the three elements — environment, operations, and the management metasystem — and the full complement of systems and channels. Here Pérez Ríos makes a deliberate, and characteristic, choice of order: study the **metasystem first (System 5, then 4, then 3), and only then the operational system (3, 3*, 2, 1)** — because in design you cannot sensibly specify operational units before you know the organization's nature and purpose (printed p. 86). System 5 is worked through checklists (is the vision written? the mission? the strategic targets?), System 4 through its adaptation apparatus (does it hold simulation models, scenarios, prospective studies? an Operations Room?), System 3 through its three duties toward the units — setting goals, negotiating resources, holding accountability — and System 1 through the horizontal and vertical relationships of each operational unit (printed pp. 87–134).

Stage 4 — Coherence across levels. Verify that the identity and purpose declared at the top are shared, suitably adapted, by every sub-organization; that the strategic plans at each level are mutually compatible; and that the channels connecting same-numbered systems at adjacent levels actually exist (printed pp. 135–139).

The move that makes it a method. Running underneath all four stages is a single, repeated test applied to *every* component of the model. For each system, channel, and homeostat, the analyst asks three questions in order: **Does it exist? Is its quality sufficient to perform its function? Does it actually perform — efficiently and effectively?** (stated at printed p. 69, applied throughout). This exists/quality/performs triple test is what converts the VSM from a diagram you nod at into a checklist you can fail. It is the most underrated thing in the book and the most portable.

What Is Actually Unique Here

The software. VSMod® is the feature that distinguishes this book from every other VSM method, and Pérez Ríos is explicit that it exists to fix the model's adoption problem (Preface; Chapter 4, printed pp. 171–200). It lets an analyst build a complete VSM, add unlimited recursion criteria and levels, and navigate them by mouse-click through three linked maps — a Global Map of the whole structure, plus a Map 1 and Map 2 that show the system-in-focus with its parent and child levels, reproducing Beer's three-level working style on screen. It renders the full detailed model of any focus — all five systems, System 3*, the homeostatic loops, transducers, amplifiers and attenuators, present and future environments — and lets the analyst* attach information in any format to any component*: text, audio, video, links, and, notably, **embedded System Dynamics simulation models** inside System 4 (printed pp. 197–198), realizing the VSM-plus-System-Dynamics integration he developed with Schwaninger. There is a collaborative version for several analysts working over the internet. No other VSM method ships with an instrument; this one is built around one.

Multi-criteria recursion. The treatment of recursion as a sphere of intersecting criteria rather than a single hierarchy is genuinely his, and it is what the software exists to make usable. Most VSM treatments quietly assume you already know how to slice the organization into levels. Pérez Ríos treats the slicing itself as an open, multi-dimensional design decision — which is why his method fits organizations, like a university or a region, that have no single natural hierarchy.

The pathology catalogue. Chapter 3 is a systematic, named taxonomy of organizational pathologies (printed pp. 141–169), modelled on medical diagnosis and organized into three families: **structural** pathologies (defects in the vertical unfolding — a missing organization at a level, absent first or middle levels, entangled memberships); **functional** pathologies, catalogued system by system (an ill-defined identity or "institutional schizophrenia" in System 5; the "headless chicken" of an absent System 4; the dissociation of the S4–S3 homeostat; an authoritarian or hypertrophied System 3; a "poor System 3*"; an authoritarian System 2; autopoietic"beasts" in System 1 that pursue only their own survival); and **information and communication** pathologies (missing information systems, absent or undersized channels, missing algedonic channels). Each carries a code and a one-line signature. It is the most usable diagnostic checklist in the VSM literature — you can walk an organization against it.

A four-stage procedure explicit enough to teach. Taken together, the staged sequence, the triple test, the matrix and the pathology list make this the VSM method a non-specialist analyst can actually follow. That is its reason for existing, and it succeeds at it.

Strengths, Best Case, Limits

Strengths. It is the most procedurally complete VSM method available: a repeatable sequence, a test that turns the model into pass/fail findings, a reusable worksheet, a named fault catalogue, and a software instrument. It handles multi-criteria recursion that other methods can only gesture at. And it is unusually strong on public-sector and territorial systems.

Best application case. The method is at its best on a large, structurally complex organization that can be legitimately cut into levels more than one way, whose environment is genuinely intricate, and where an analyst or small team will do sustained expert work over time — public universities, regional and city governments, health systems, multi-scale territorial bodies. It rewards the analyst who wants rigor and documentation over speed, and it is the right choice when the recursion structure is itself in question. It is also the natural method to pair with quantitative modelling, because the software will hold System Dynamics models inside the sensing function.

Limits. Three are worth naming. First, **software dependence**: the multi-criteria recursion and the matrix's full power are pitched as needing VSMod®, the book is tied to a 2007-era version, and the durability of any method that leans on one program at one university's web address is a real risk. Second, **it is expert-led desk work**: the analysis is driven by an analyst or consultant through interviews, document review and checklists, not by the organization's members; participation is a separate, optional instrument (Team Syntegrity, Chapter 5), not part of the diagnostic method. Third, the very exhaustiveness that is its strength makes it **heavy** — the completeness that suits a university or a region can be disproportionate for a small firm that needs a fast structural read, for which Beer's four charts or Pfiffner's five-minute napkin are better fitted.

Part I — The Political Environment the Method Was Built For

The word "Sustainable" is in the title, and it is the single most important thing to read correctly, because it does not mean what a reader coming from Espinosa and Walker will expect.

For Pérez Ríos, *sustainable means viable*. The concluding chapter is explicit: the aim is "always to ensure that both of these are **viable**, that is, that they have the capacity to maintain a separate and long-lasting existence regardless of the changes... in the environment" (printed p. 218). Throughout the book "sustainable" functions as a synonym for self-sustaining or long-lasting — a structural property of an organization that keeps its identity through change — not as a claim about ecology, climate, or the biosphere. Ecological problems appear, but as *examples of structural pathology* (cross-border pollution as a case of a missing organization at a recursion level)

or as candidate topics for a group method, never as a substantive politics. This is the sharpest single contrast in the whole set of methods: Espinosa and Walker put ecological-societal sustainability at the centre and build their politics around it; Pérez Ríos uses the same word to mean structural viability and builds no such politics at all.

That is not a criticism — it is a clarification of genre. Pérez Ríos's book is a *technical* text in the most precise sense: it treats the design of viable structure as an engineering problem that can be specified, taught, and tooled, largely independent of whose interests the structure serves. Like Pfiffner and like Beer, he makes the technocratic move of bracketing the question of interest. He makes it more completely than either, because his instrument — a piece of analysis software — has no place in it to represent a contest of interests at all.

Where the method does have an implicit political home, it is the **public sector**. The flagship case is a public university and its region; the structural-pathology examples are dominated by public problems — supra-municipal transport, water and waste; metropolitan areas without a matching authority; trans-provincial rivers; international courts; cross-national ecology — and his System 5 examples include an Autonomous Community's Economic and Social Council and a University's Social Council. The method is pitched, more than any other in this set, at government, regional, city and university scales. But it approaches them as problems of *structural viability* — is there an organization matched to each relevant environment, and does it hold the five systems? — not as problems of democratic contestation. A region with a missing middle recursion level is diagnosed the way a body with a missing organ is diagnosed: as a functional gap, not as a political settlement.

The best political environment for this method, then, is a public or quasi-public institution that genuinely wants a rigorous structural diagnosis and has an analyst to do it — a university, a regional administration, a health service — and whose leadership will accept findings expressed as structural faults rather than as questions of power. Where that holds, the method's political neutrality is an asset: it gives a shared, checkable language across contending departments. Where the real problem *is* a contest of interests — who should govern a metropolitan area, whose region absorbs a cost — the method will correctly locate the missing structure and will, by design, have nothing to say about who should fill it.

The one place the book's register changes is Chapter 5, on Team Syntegrity — Beer's non-hierarchical group method, thirty participants and twelve topics on the frame of an icosahedron. This is the only part of the book where "democratic", "non-hierarchical", "participatory" and even "future generations" appear, and where social, economic, political and ecological issues are named as legitimate subjects. It is genuinely participative — but it is a separate, optional instrument, recommended mainly to make the System 4–System 3 conversation work, not a property of the diagnostic method itself.

Part II — The Internal Power Constellation It Presupposes

Inside the organization, the method's assumption about power is the same as Beer's in kind but not in tone: the pen is held by an expert. Where Beer wrote for the manager working alone and Pfiffner for a confidential core team, Pérez Ríos writes for the **analyst or consultant** working through interviews, documents, checklists and software — an even more expert-centred constellation, because the instrument itself presumes a trained operator.

The analyst drives; the organization is the object of study. The four stages are performed by someone applying the exists/quality/performs test to each component and recording the result in VSMod®. The members of the organization supply information; they do not co-construct the model. This is the opposite pole from Espinosa and Walker, where the whole point is that the community builds the model itself, and it is a legitimate choice — expert diagnosis produces rigor and comparability that participative construction does not — but it is a choice about who holds authority over the organization's self-description, and it should be made knowingly.

Participation is real but quarantined. Democracy, non-hierarchy and participation are not absent from the book; they are concentrated entirely in the Team Syntegrity chapter. The diagnostic and design method itself contains no theory of stakeholder or employee legitimacy. Employees appear as objects — units to be granted autonomy and audited, not authors of the analysis. Works councils, unions and labour appear only incidentally, as examples of a communication homeostat. This is a clean separation: the method is the analyst's; participation is a distinct tool you may bolt on. A reader who wants the diagnosis to double as a legitimacy-building process will have to supply that themselves, most naturally by running the analysis and the Syntegration together.

Where the internal preconditions fail. They fail where there is no analyst with the time and training to do sustained expert work; where the organization will not accept a structural diagnosis it did not itself build; where the software is unavailable or its multi-criteria complexity is overkill; and where the live problem is a contest of interests that a structural diagnosis will name but not resolve. In those settings the *content* remains valuable — the pathology catalogue and the triple test are usable by anyone, on paper — but the method as designed, instrument and all, needs an expert to hold it.

Placed beside the others, the pattern completes cleanly. The four methods answer the same question — *who holds the pen, and with what?* — in four ways:

- **Beer:** the manager, alone, with four hand-drawn charts.

- **Pfiffner:** a confidential core team of three to five, with a variety-measured seven-step process.
- **Espinosa & Walker:** the whole community, in workshops, building the model together.
- **Pérez Ríos:** the analyst, with a checklist, a matrix, a pathology catalogue, and a piece of software.

They are not rival readings of the model. They are four different placements of authority over the model, each right for a different organization.

What to Borrow

The exists / quality / performs triple test. This is the most portable thing in the book, and it needs no software. For any component of any VSM analysis, ask in order: does it exist, is it good enough to do its job, and does it actually work? It converts a diagram everyone agrees with into findings some things fail.

The Recursion Levels–Critical Factors Matrix. Even on paper, rows-as-levels against columns-as-critical-factors forces the vertical analysis to be complete. It is worth building for any multi-level organization before drawing a single system.

Multiple recursion criteria, made explicit. Before you unfold an organization, ask how *else* it could legitimately be unfolded — by territory, by product, by function — and decide the cut consciously. Pérez Ríos's insistence that this is an open, multi-dimensional choice prevents the common error of treating the existing org chart as the only possible recursion.

The pathology catalogue as a diagnostic checklist. Walk a struggling organization against his named faults — missing recursion level, ill-defined identity, headless-chicken System 4, hypertrophied System 3, poor System 3*, authoritarian System 2, autopoietic units. It is the fastest structured read available, and it names things leaders recognize.

What does not transfer for free is the instrument. VSMod® is what makes the multi-criteria, multi-level analysis tractable at full scale; without it, the method reverts to the same paper limits as the others, and its distinctive advantage narrows to the matrix and the pathology list. Adopt the discipline knowing that the machine is part of the method.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I;

ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5. Software: VSMod®, University of Valladolid (www.vsmod.org).

Chapter 5

Naming the System

Espejo and Reyes's VIPLAN Method — Diagnosis That Begins by Deciding What the Organization Is

In Brief

Every VSM analysis has to answer one question before any other: *what is this organization, exactly?* Beer's handbook assumes you can already see the system-in-focus. Raúl Espejo — who ran the operations of Beer's Cybersyn project in Allende's Chile, and spent the four decades since applying the model in Colombia, Britain and Sweden — treats that assumption as the whole problem. His method, VIPLAN, is built on the claim that naming the organization is not a preliminary to the analysis; it *is* the first and most consequential act of it, because the same set of activities can be two entirely different organizations depending on whose viewpoint names them, and each naming has structural consequences that follow with unusual directness.

Organizational Systems (2011), written with Alfonso Reyes, is the definitive statement of VIPLAN. It supplies what Espejo says Beer's *Diagnosing the System* lacked — "a guide to apply the VSM but not an epistemology" (Preface). Onto Beer's model it grafts two things: a second-order-cybernetic epistemology (drawn from Maturana and von Foerster) that makes identity a matter of *the observer's* judgment, and a step-by-step method — TASCOI, the unfolding of complexity, the recursion/functions table, and the design of structural mechanisms — with dozens of real, named cases.

This chapter describes that method, names what is unique about it, identifies its best application case, and states its limits. It then makes explicit the politics the method carries — which, unusually for a VSM book, is worn on the surface: an ethics of autonomy taken from von Foerster, a reading of hierarchy as "bad cybernetics", and a body of work rooted in the public sector and the state, from a national audit office to a country's nuclear-waste settlement.

A note on terminology, because the field is loose about it. This book distinguishes three things it is easy to conflate: the **Viable System Model** (Beer's model — the object); the **VIPLAN method** (the diagnostic and design procedure of Part II); and the **VIPLAN methodology** (the wider systemic problem-solving approach of Part III). It does not use the phrase "the VIPLAN model." This chapter follows the book.

Two Prisons in One Building

Espejo asks you to write the identity of a prison. From the governor's chair, the transformation is clear enough: the prison receives people convicted by a criminal court and holds them as inmates, in order to protect the community. The community is the customer. The courts are the suppliers. The staff are the actors. Parliament and the justice system are the interveners who set the rules. A governing body owns the transformation (printed pp. 127–128).

Now write it from an experienced inmate's chair. The transformation is *improving one's criminal technique*. The actors are the seasoned inmates. The customers — the ones the transformation is performed upon and for — are the newer prisoners. Same building, same people, same daily routine; a completely different organization, with a completely different structure implied.

The point is not that one naming is true and the other false. It is that *both are operating in the same building*, and the governor's structure — which treats inmates as inputs to be processed rather than as customers of anything — has no place in it for the second transformation, which therefore runs unmanaged and unseen. Espejo's method makes this visible on purpose, because the act of naming determines what the subsequent structural analysis can and cannot see. Get the naming wrong and every diagram after it inherits the error.

This is why VIPLAN begins where it does. Beer's model tells you what a viable system must contain; it does not tell you which organization you are looking at, or from whose standpoint. VIPLAN makes that standpoint an explicit, defended, revisable choice.

The Method: Five Steps, Two Modes

The VIPLAN method runs in either of two modes — **Mode I, diagnosis** of an existing organization, or **Mode II, design** of a new or reconfigured one — and proceeds through five steps (Chapters 7–9). The steps below are described in diagnosis mode; in design mode the same instruments run prospectively, the distinctions being chosen rather than found.

Step 1 — Naming systems (TASCOI). Establish the organization's identity through a black-box description of what it *does*, not what it claims. The canonical form is Checkland's: *the system does X by means of Y in order to achieve Z*. The naming is then disciplined by the **TASCOI** checklist (printed pp. 124–126):

- **T**ransformation — the inputs turned into outputs;
- **A**ctors — those who perform the transformation;
- **S**uppliers — of the inputs;
- **C**ustomers — who receive the outputs;

- **O**wners — those with an overview of the whole transformation, accountable for it;
- **I**nterveners — those who set the parameters within which it operates.

Espejo draws a sharp line between TASCOI and Checkland's better-known CATWOE. TASCOI names "meaningful chunks of complexity *in the world*" — real actors, real suppliers, a human communication system that exists whether or not anyone models it. CATWOE names *ideas about* the world, mental constructs for reflection (footnote, printed p. 126). This distinction — the *operational domain* where the organization actually happens versus the *informational domain* where people reflect on it — is load-bearing for the entire method, and it is Espejo's chief epistemological contribution.

The naming is not done by an expert alone. Espejo introduces the **enabling viewpoint**: the person or group using the method produces an initial identity *hypothesis*, then facilitates iterative identity workshops with participants until the names stabilize. Who *ought* to count as a customer or an actor is, he insists, a values question and a boundary judgment, not a matter of logic (printed p. 128).

Step 2 — Technological and structural modelling. Model the transformation itself — the actual technological chain by which inputs become outputs — and the structural options that chain permits, as the ground on which structure will be built. Different production technologies allow the work to be divided in different ways, so this step surfaces the structural possibilities among which the next step will choose. The book treats it far more briefly than the steps on either side; it functions mainly as a bridge from the naming to the unfolding.

Step 3 — The unfolding of complexity. This is Espejo's signature move, and the term is his (he credits himself, Espejo 1989). An organization is decomposed into nested viable systems by applying **complexity drivers** — the criteria along which its variety naturally divides. He names four: **technology, customer/supplier, geography, and time** (printed p. 138). Each driver yields a different way of cutting the organization into levels, and the analyst chooses and combines them to produce a *structural recursion* — viable systems within viable systems. Crucially, he distinguishes **chunking** (dividing into genuinely autonomous parts, which amplifies the whole's capacity to respond) from **fragmenting** (dividing into dependent parts that must refer everything upward). And he is emphatic that the unfolding is *not* a management hierarchy and *not* an org chart — it maps relations of autonomy and operation, not power (printed p. 161). A primary activity is one that produces the organization's transformation and must therefore develop real autonomy; a regulatory or support activity serves the primary ones and must integrate into them rather than acquire autonomy of its own — and which is which depends entirely on the organization's purpose (a university library is regulatory; a public library is primary).

Step 4 — Distributing discretion. Here the method reframes the tired centralize-versus-decentralize argument as a design question about **discretion** — the explicit agreement that particular resources are a particular role's responsibility. The tool is the **Recursion/Functions Table**: primary activities as rows (preserving the unfolding), regulatory functions as columns, a mark in a cell where that activity holds discretion for that function. Read the columns and you see which functions are centralized; read the rows and you see each unit's real autonomy. Espejo gives **five criteria** for decentralizing a function to a primary activity: it is a critical success factor for that unit; its application is particular to that unit; demand for it within the unit is high; the resources it needs are available within the unit; and those resources are distributable (printed p. 167). If all five hold, decentralize.

Step 5 — Designing structural mechanisms. For every function, design the **cohesion mechanism** that lets a unit be autonomous without the whole flying apart. It has four components (printed, Chapter 9): a channel issuing the non-negotiable centralized rules (the policies that reflect the organization's ethos, values and legal obligations); a two-way **resource-bargaining** channel carrying committed outcomes; a **monitoring** channel confirming the commitments hold (of which auditing is one instance); and a **coordination** function that pre-empts problems between units. The **adaptation mechanism** — connecting cohesion, intelligence and policy — is designed the same way. His cautionary tale is the organization whose thirty-two regions each diligently decentralized their information systems, ending with thirty-two incompatible systems and no consolidated view of anything, because no one had designed the cohesion mechanism (printed, Chapter 9).

What Is Actually Unique Here

TASCOI and the primacy of naming. No other VSM method makes identity a formal, contested, first-class step. TASCOI is arguably the most widely borrowed single tool in the applied-VSM literature, and this is its home.

A second-order-cybernetic epistemology. Espejo grounds the whole method in the claim that a system is a *distinction drawn by an observer* — "everything said is said by an observer", after Maturana — and that organizations are networks of recurrent communication that achieve *closure*, quoting Beer: "Closure is the talisman of identity" (printed, Chapters 1, 5). Closure here means the organization's self-referential pattern of relations — the network of communications that keeps producing itself, so the organization stays the same system while everything in it changes. This is why identity can be named two ways in one prison, and why the method treats structure as *relations between roles*, explicitly not the formal hierarchy or the org chart. It is the most philosophically serious of the VSM methods, and the epistemology is not decoration — it is what forces the naming step to be taken seriously.

Residual variety. Espejo's refinement of Ashby's Law is that a manager does not need variety to match *all* the states of what they manage — only the **residual variety** left over after the managed unit's own self-regulation has absorbed what it can. Strengthen local self-regulation and coordination, and less residual variety travels upward, which means leaner management and greater local autonomy at the same time (printed pp. 63–70). It is the cleanest available argument for why devolving is not a loss of control but a redesign of where variety is absorbed.

The recursion/functions table and the five decentralization criteria. These turn "how much autonomy?" from a values fight into a function-by-function analysis with a testable rule. Alongside Pérez Ríos's critical-factors matrix, it is one of the two most useful worksheets in the applied literature.

Strengths, Best Case, Limits

Strengths. VIPLAN is the most epistemologically grounded VSM method in the set, and among the most procedurally explicit — a rare combination. It has an unusually deep bench of real, named, public cases. And it is genuinely participative in a way most expert-led methods are not: the naming and discretion workshops are constitutive, not cosmetic.

Best application case. The method is at its best on an organization whose very identity is contested or unclear — a merged entity, a public body pulled between mandates, a regulator, an organization that has grown past its own self-description — and where a skilled facilitator can run identity and discretion workshops with real stakeholders over time. It is the method to reach for when the first honest question is not "is the structure right?" but "what is this organization, and for whom?" Its natural home is the public and quasi-public sector, where that question is hardest and most consequential.

Limits. Three. First, it is **heuristic and non-algorithmic** by the authors' own admission — "in the end the methodology is fundamentally heuristic" — and its quality depends heavily on the skill of the enabling viewpoint; a weak facilitator produces a weak naming and everything downstream inherits it. Second, building the full VSM from the recursion/functions table is, in the book, handed off to the VIPLAN software rather than fully worked on the page. Third, the epistemology that is its great strength is also a barrier to entry: managers who want a fast structural read will find the Maturana and the operational-versus-informational-domain apparatus heavy going, and will be better served by Hoverstadt or Beer's charts.

Part I — The Political Environment the Method Was Built For

Of all the VSM methods, VIPLAN is the one whose politics is explicit rather than bracketed — and it comes from the same place the method itself does: Chile.

Espejo states it plainly in the Preface. In 1971–73 he was, on his own account, deeply involved in **Project Cybersyn**, Beer's attempt to build a real-time regulatory system for the whole economy under Salvador Allende — Beer's "third way", neither centralist planning nor laissez-faire. And his retrospective verdict on it is candid and formative: the approach was "bold but too optimistic; producing effective relationships between stakeholders and policy-makers was far more complex than building up communication networks and information systems" (Preface). That lesson — that the hard part is *relationships*, not information plumbing — is the stated reason the next great cybernetic project of his career, the work with Colombia's National Audit Office, put stakeholder relationships at its centre. VIPLAN is, in a real sense, the method Espejo built to correct the thing Cybersyn got wrong.

The ethics is drawn directly from second-order cybernetics, and it is stated as an imperative: after von Foerster, "ethical behaviour is one that increases possibilities for others and in particular does not restrict their viability" (printed, Chapter 1). Systems thinking is defended on ethical grounds — as the discipline that "helps us to connect distant events and work out the hidden consequences of our actions." And hierarchy is read critically throughout: uneven power concentrates "in the hands of the few most of the influence to produce knowledge"; hierarchies are, in his phrase, "bad cybernetics" (printed, Chapter 6). The normative core of the whole method is the complementarity of **autonomy and cohesion** — autonomy understood etymologically as governing oneself *and* cybernetically as producing one's own products, held in balance with the cohesion the whole requires.

Democracy is treated explicitly where the cases demand it. In the Swedish nuclear-waste study, Espejo writes that "in democratic societies it is unacceptable for central governments to impose unilateral solutions", and that affected communities "expect participation in the appreciative process from the beginning", concerned above all with "the legitimacy of decision processes" (printed, Chapter 11). This is not incidental colour; it is where the method's participative machinery earns its keep.

And the application domain is, more than any other method in this set, the **state and the public sphere**. The anchor project is the National Audit Office of Colombia; the worked cases include the City of Bogotá's audit office, Colombia's National Registry, a national higher-education examinations body, prisons, and the Swedish nuclear-waste settlement — alongside private firms like NatWest, Satena and a pension-fund company. Espejo claims applicability "from large private companies and public institutions

to small companies and NGOs in many different countries", and the claim is backed by cases on four decades and three continents.

The best political environment for VIPLAN, then, is a setting where identity is genuinely contested and legitimacy genuinely matters — a public institution, a multi-stakeholder settlement, a merged or mandate-torn body — and where those in charge will accept that naming the organization is a values question to be worked out with stakeholders, not an answer to be handed down. Where that holds, the method's explicit ethics is an asset. Where a commissioner wants a structural fix without reopening the question of whom the organization is for, VIPLAN will keep insisting on the question they wanted to skip.

Part II — The Internal Power Constellation It Presupposes

VIPLAN sits at a distinct point on the spectrum of who holds the pen — neither the solo diagnostician nor the whole self-governing community, but something in between that the book names precisely: the **enabling viewpoint**.

An expert catalyzes; participants decide the naming. The enabling viewpoint — an analyst, consultant or internal catalyst — seeds the initial identity hypothesis and designs the process, but the naming itself is settled in workshops with the relevant participants, and the discretion analysis is negotiated with the roles that hold the resources. Espejo is explicit that expert-only work fails at exactly the point that matters: "it is not the analytical skills of a few people creating sophisticated models that will ground these meanings" (printed, Chapter 11). The pen is shared, but someone is responsible for holding it out.

The method presupposes stakeholders who can be brought into the room. This is its central internal precondition. TASCOI's two prisons only become useful if the people who hold the divergent namings can be convened and their differences worked through; the discretion workshops only work if the roles negotiating their own boundaries actually attend. Where power is too concentrated or too defended for that convening to be honest — where the naming is effectively dictated — the participative machinery idles and the method degrades into an expert diagnosis wearing a participative costume.

Structure is relations, not the org chart — which redistributes authority over the self-description. Because Espejo defines structure as the relations between roles rather than the boxes on the chart, the analysis routinely reveals that the real organization is not the official one, and that a primary activity everyone treats as central is regulatory, or an unrecognized transformation is running unmanaged (his National Registry case turns on exactly such an unrecognized fourth primary activity). Naming these things reallocates status and attention. If that reallocation is off the

table, the naming cannot be done honestly, and Step 1 — the step the whole method rests on — is compromised.

Where the internal preconditions fail. They fail where identity is not permitted to be an open question; where stakeholders with divergent namings cannot be safely convened; where the organization wants a structural tune-up without reopening who it is for; and where no facilitator of sufficient skill is available to hold the naming. In those settings the *tools* remain usable — TASCOI as a solo desk exercise, the recursion/functions table, the residual-variety lens — but the method as designed, participative and identity-first, needs a room it can actually convene.

Placed beside the others, VIPLAN occupies the middle of the range:

- **Beer:** the manager, alone, with four hand-drawn charts.
- **Pfiffner:** a confidential core team of three to five, with a variety-measured seven-step process.
- **Pérez Ríos:** the analyst, with a checklist, a matrix, a pathology catalogue, and software.
- **Espejo & Reyes:** an enabling viewpoint that seeds the analysis and then convenes the stakeholders to name the organization and negotiate its discretion.
- **Espinosa & Walker:** the whole community, building the model together over months.

They are five placements of authority over the model. VIPLAN's distinctive claim is that the placement question and the *naming* question are the same question — because who gets to say what the organization is, is the first and deepest distribution of power in it.

What to Borrow

TASCOI, always. Before any structural work, name the organization: transformation, actors, suppliers, customers, owners, interveners — and do it from more than one viewpoint. The divergences are not noise; they are the diagnosis. This single tool transfers into any of the other methods and improves all of them.

The operational-versus-informational-domain test. For any proposed change, ask whether it is embodied in the actual relations and resources of the organization (operational) or whether it lives only in a document and a set of good intentions (informational). Espejo's warning is that unembodied reflection "remains outside the organization's boundaries" — a plan that changes no relation changes nothing.

Residual variety. When someone argues that devolving control means losing it, reframe: management needs to match only the variety a unit's

own self-regulation cannot absorb. Strengthen local regulation and coordination and the control problem shrinks at the top *and* autonomy grows at the edge — the same move, not a trade-off.

The recursion/functions table with the five criteria. For any multi-unit organization, lay activities against functions and decide each cell against the five decentralization criteria. It converts the perennial centralize/decentralize argument into a function-by-function analysis that can actually be settled.

What does not transfer for free is the room. VIPLAN's power comes from convening the people who hold the competing namings and letting them work it through. Borrow the tools solo if you must, but the method's real yield — an identity the organization actually shares — needs the participation the tools were built to structure.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.

Chapter 6

The Three Clocks

Schwaninger's Model of Systemic Control — Why the VSM Needs a Second Axis

In Brief

The other methods in this set apply the Viable System Model. Markus Schwaninger, of the University of St. Gallen, argues that the VSM — for all its power — answers only one of the two questions an organization must get right, and that using it alone quietly leaves the other unanswered. The VSM tells you how variety should be *distributed* across levels of recursion. It does not, on its own, tell you what a management team should be *steering by* — and Schwaninger's claim is that any organization is simultaneously governed by three sets of control variables that belong to different logical levels and routinely contradict each other. Reconciling them, he writes, "is a core task of leadership."

Intelligent Organizations (2nd ed., 2009) is therefore not a VSM manual but an integration. It sets the VSM (for structure) alongside four other models — the **Model of Systemic Control** or MSC (for activities), Team Syntegrity (for behaviour), normative-management theory (for ethos and identity), and System Dynamics (for time) — into a single framework for designing what he calls organizational intelligence. The spine of that framework is the MSC and its **three logical levels of management: operative, strategic, and normative**, an axis that runs *orthogonal* to the VSM's recursion.

This chapter describes that framework, isolates what is genuinely unique in it — the two-axis scheme, the logic of *pre-control*, and the idea of viability *beyond* survival — identifies its best application case, and states its limits, including its frank abstraction. It then locates the politics, which in Schwaninger's system is not bracketed but built in: the normative level explicitly carries *legitimacy*, the ethical and even aesthetic dimension is placed on the management agenda, and Team Syntegrity is offered, in his words, as "an architecture for democracy."

A note on the source. Schwaninger states that he is not writing "a book about how to run a company" but one "about how to look at the world differently" — a heuristic framework, "not a quick fix." Read as such: this is the most theoretically ambitious of the VSM methods, and the least like a checklist.

The Bank That Was Perfectly Managed — at the Wrong Clock

Consider — as a composite drawn with hindsight, not a named case — a bank in the mid-2000s that was, by the only measures its leadership watched, superbly run. Profit was at record levels. Liquidity was ample. Every operative dial read green. Two or three years later it was insolvent, rescued or absorbed, its name a byword for the crisis.

Schwaninger's diagnosis of this pattern — and he writes it directly against the 2007–08 banking collapse (printed pp. 71–74) — is not that the managers were greedy or stupid, though he quotes a finance minister to the effect that "greed has switched off the brains." It is that they were steering by the wrong clock. Profit and liquidity are the control variables of *operative* management, and they run fast. But beneath them, moving more slowly and largely invisible on a quarterly report, sit the control variables of *strategic* management — the bank's stock of value potentials, its real capabilities and the trust of its market — and beneath *those*, slower still, the variables of *normative* management: the organization's viability and its legitimacy. The bank that maximizes the fast clock while running down the slow ones looks excellent right up to the moment it is not. Its operative success was *consuming* the strategic and normative potentials that made the success possible.

This is the intuition the whole framework formalizes. An organization is governed by three clocks at once, ticking at different rates, and a management system that can read only the fastest of them is not simple — it is blind.

The Framework: Five Models, Three Levels, Two Axes

Schwaninger builds the design of an "intelligent organization" from five dimensions, each governed by a dedicated model (printed, Chapters 4–9):

Dimension	What it concerns	Governing model
Activities	what the organization does	**Model of Systemic Control (MSC)**
Structure	preconditions for effective action	**Viable System Model** (Beer)
Behaviour	the control of cognition and emotion	**Team Syntegrity** (Beer)
Ethos & Identity	the basic parameters	normative-management theory (St. Gallen)
Time	organizational dynamics	**System Dynamics** (Forrester)

The **Model of Systemic Control** is the heart of it, and its content is the specification of *which control variables* an organization needs at each of three logical levels (printed pp. 48–71):

Operative management steers by **profit and liquidity** — or, in the generalized version he builds for public bodies and social ventures, by customer, financial, social and ecological benefits. Its criterion is **efficiency**: doing things right. Fastest clock.

Strategic management steers by **existing and new value potentials** — the resources, capabilities and core competencies that must already be in place for value to be created at all. Its criterion is **effectiveness**: doing the right things. Middle clock.

Normative management steers by **viability and development**. Its orientators are the system's ethos, identity, structure, culture and dynamics; its criterion is **legitimacy** — "the ability to fulfil the legitimate claims of all relevant stakeholders." Slowest clock, longest horizon.

Binding the three together is the mechanism Schwaninger calls **precontrol** (after Gälweiler): "the creation of preconditions… at a higher logical level, which largely predetermine what can be achieved… by the lower logical levels" (printed p. 49). Viability and development set the ceiling on value potentials; value potentials set the ceiling on profit; profit sets the ceiling on liquidity. You cannot buy back a squandered value potential with next quarter's earnings — the chain runs one way, which is exactly why the well-managed bank could not save itself once the slow clocks had run down. (There is a feedback path upward too: lower-level outcomes enable or constrain the higher referents.)

The second axis. Here is the move that makes this more than a restatement of the St. Gallen management tradition, and the reason a VSM practitioner should care. The three logical levels are **not** the same as the VSM's levels of recursion, and Schwaninger insists on the distinction repeatedly (printed pp. 50, 89). Recursion is the *nesting* axis — autonomous wholes within autonomous units. Logical level is a *management-logic* axis that is present, in full, inside every one of those units. Mapped onto the VSM: Systems 1–2–3 do the operative work; System 3 in interaction with System 4 does the strategic work; System 5 does the normative work (printed p. 85). But — and this is the payoff — *all three logical levels must be exercised at every level of recursion*. It is not true that "vision is the entrepreneur's job" or "strategy is the board's": each subsidiary, each business unit, each team must speak all three languages in its own metasystem (printed pp. 90–91). Schwaninger calls this "the structuralist foundation for humanistic postulates such as autonomy, meaningful work, participation and self-fulfilment." The two axes together — three logical levels crossed with N levels of recursion — are the framework's signature, and they resolve a confusion that flattens most single-axis VSM work.

The Method: Model-Based Management, Dimension by Dimension

Schwaninger is candid that this is a heuristic, not a lock-step procedure. What follows is accordingly a lens-set rather than a procedure — an account of what looking through each lens involves, not steps to execute. But there is a method, and its governing principle is the **Conant–Ashby theorem**: every good regulator of a system must be a model of that system — from which he draws the sharp corollary that "the result of an organizational process cannot be better than the model on which the management of that process is based, except by chance" (printed p. 20). Management *is* model-building, and the framework's job is to make sure the models are good ones. Diagnosis and design proceed dimension by dimension:

For **activities**, assess the control variables at each logical level and the integrity of the pre-control chain; use configuration analysis to compare a firm's business-system configuration against rivals', and business design to move from the status quo through an ideal design to a plan.

For **structure**, run a VSM diagnosis — are Systems 1–5 present, adequately resourced, and properly connected at each recursion level? — then redesign.

For **behaviour**, run *syntegrations*: the fully specified Team Syntegrity protocol (thirty people, twelve topics on the edges and vertices of an icosahedron, iterated until the great part of the relevant information is shared across the group), which he presents as the means of implementing the metasystem's Systems 3–4–5 conversation.

For **ethos and identity**, work discursively — ask "who are we?" and, in his own preferred move, "who are we *not*?", because a negative preclusion absorbs more variety than a positive statement — and assess whether the ethos is healthy or pathological, conscious or unexamined.

For **time**, build System Dynamics simulations to test whether interventions across dimensions actually combine well. His own finding, which he wryly calls "Markus' Paradox," is that combining interventions is not in itself superior to punctual ones; superiority appears only with *synchronization* — from which he draws the capstone maxim that "the basis of effective time management is starting earlier, not acting faster."

What Is Actually Unique Here

The two orthogonal axes. No other VSM method cleanly separates *what you steer by* (logical level) from *where the variety sits* (recursion), and the separation dissolves a genuine confusion — the endless argument about whether strategy "belongs" at the top is revealed as a category error, because strategic management is a logical level present at every recursion, not a floor of the building.

Pre-control. The insistence that the slow variables predetermine the fast ones — viability before value potential before profit before liquidity — is the clearest available account of why organizations that optimize their

reported numbers can destroy themselves, and why intervention has to start "earlier, not faster."

Viability beyond survival. Schwaninger refuses to let viability be the ceiling. Mere viability, he warns, risks "pathological autopoiesis" — self-perpetuation that harms the larger whole. Above it he places *development* (after Ackoff): "a system's growing ability and desire to fulfil its own and others' needs and legitimate desires", which can require the deliberate transformation of identity itself (his examples: a sawmill becoming Nokia, which began as a pulp mill; Preussag becoming a travel company). This is a normative ambition most VSM writing does not reach for.

Integration and model-based management. The framework is unique in binding four cybernetic and systemic models into one, and in insisting — via Conant–Ashby — that qualitative structure (VSM) and quantitative simulation (System Dynamics) are two halves of a single discipline. The MSC exists partly to give simulation models the pre-control skeleton they otherwise lack.

Strengths, Best Case, Limits

Strengths. It is the most conceptually complete of the VSM methods — the one that connects structure to strategy to finance to ethics to time in a single scheme. It carries a genuine normative and ethical dimension most management models omit. And it disciplines the loose talk of "levels" that bedevils single-axis VSM practice.

Best application case. The framework is at its best with the top team of a medium-to-large organization at an inflection point — one whose operative numbers look fine but whose leadership senses that its value potentials or its legitimacy are eroding, and who are willing to invest in a serious, model-based diagnosis rather than a quick structural fix. It suits organizations facing a possible transformation of identity (a mature firm whose market is dissolving), and it pairs naturally with simulation where the dynamics are genuinely counterintuitive. It is the method to reach for when the presenting problem is *"we are winning on the metrics and losing somehow anyway."*

Limits. Schwaninger says so himself: this is "not a quick fix", "not a book about how to run a company", and parts of it are "work in progress" with "some empirical tests still pending"; his own integrative simulation model he calls "too broad and too abstract" to validate on data. The practical how-to lives partly in appendices and in cross-referenced earlier work rather than in step-by-step form on the page. And there is a real tension: a framework built by "one rational designer" that also espouses distributed, participative, democratic ideals — the very critique once leveled at the VSM, now internal to the method. The book's answer to it is Team Syntegrity; how completely that answer resolves the tension is arguable. A manager who wants a fast, hands-on structural read will find Beer's charts or Hoverstadt's

symptom catalogue far more immediately usable; Schwaninger rewards the reader who wants depth over dispatch.

Part I — The Political Environment the Method Was Built For

Where most VSM methods bracket the question of *whose interests the structure serves*, Schwaninger builds it into the model — at the normative level, which is precisely the level single-axis VSM work tends to underplay.

The normative criterion is **legitimacy**, and he defines it expansively: "the ability to fulfil the legitimate claims of *all* relevant stakeholders", explicitly including "external constituents such as the community, government and future generations" (printed p. 140). Future generations are named as stakeholders, on the authority of Hans Jonas's ethics of responsibility (printed p. 141). This is not a footnote; the entire eighth chapter places ethics — and, unusually, aesthetics — on the management agenda as first-class normative concerns, and criticizes their routine omission. Schwaninger endorses the procedural position of the St. Gallen ethicists (that ends and values "must be submitted to institutionalized processes of ethical reflection and democratic legitimation") and then goes further than most management writers dare, venturing substantive value content — cooperativeness, trust as social capital, even a catalogue of virtues — while explicitly refusing relativism ("if anything goes… nothing goes any more") without lapsing into a single fixed value system.

His sharpest political-economic edge is the critique of **short-termism** (printed pp. 71–74): an extended indictment of profit-maximization and financialization, illustrated through the scandals of the 2000s and the banking crisis, with concrete institutional remedies (long-term incentive design, blocking periods on options, changes to the shareholder mix). And his positive exemplars carry a politics of their own — the Mondragón cooperative federation, praised for democratic collective ownership that enables "joint sense-making"; restorative ecological management; the triple bottom line built into the operative level of the generalized MSC.

Participation is treated seriously, and critically. It is "a lever for complexity absorption" and necessary — but "not sufficient", and it can be "ill-conceived". Its proper institutional form, for Schwaninger, is **Team Syntegrity**, which he frames flatly as "the design for democratic management" and "an architecture for democracy": a heterarchical structure in which influence is distributed equally across thirty members and twelve teams, engineered to defend against both the domineering individual and groupthink.

The application domain is deliberately universal — the generalized MSC is built for "private firms, public institutions, non-governmental organizations and social initiatives" — and the exemplars range from multinationals

through EU-funded regional systems and a national transport ministry to universities, a hospital, and Mondragón.

The best political environment for this framework, then, is an organization whose leadership is willing to treat legitimacy and long-term viability as real objects of management, not public-relations afterthoughts — and to accept that the normative level constrains the operative one, that some profitable things are illegitimate and some legitimate things unprofitable, and that the two must be balanced rather than collapsed. Where that willingness exists, Schwaninger gives it the most developed apparatus in the VSM literature. Where a commissioner wants only the fast clock optimized, the framework will keep insisting on the two slower ones they were trying not to look at.

Part II — The Internal Power Constellation It Presupposes

Schwaninger's method sits at an unusual double position: it is **expert-and-management-model-led in its execution**, and **distributed-participative in its ideal** — and the tension between the two is real and openly present in the book.

Balancing the three clocks is leadership's job. The reconciliation of contradictory control variables across the three logical levels is named "a core task of leadership"; the framework is addressed to managers, executives and advanced students, and its diagnostic apparatus is wielded by someone who can build and read the models. In execution, the pen is held by the top team and its advisers.

But management is a distributed function, by doctrine. Schwaninger's own commitment — after von Foerster, "each participant is also a manager" — is that managerial functions belong across the whole system, and that all three logical levels must be exercised at every recursion. The framework thus *prescribes* that vision and strategy not be hoarded at the top even while it is *deployed* largely from the top. The resolution he offers is Team Syntegrity: a concrete, heterarchical, equal-influence mechanism through which the metasystem's strategic and normative conversations can actually be conducted by many rather than dictated by few.

The precondition is a leadership that will submit its own steering to reflection. The method presupposes decision-makers willing to make their models explicit and testable (Conant–Ashby applied to themselves), to ask "who are we *not*?" honestly, and to accept that the normative level sits above the operative one in the pre-control chain — which means accepting constraints on profit-seeking that a purely operative leadership would never entertain. Where leadership will not submit to that reflection, the framework has no lever; it is not a method that can be imposed on an unwilling metasystem, because the metasystem *is* the thing being asked to change.

Where the internal preconditions fail. They fail where leadership treats the normative level as decoration; where the appetite is for a structural tune-up rather than a reflection on identity and legitimacy; where there is no capacity or patience for model-based work; and where participation is unwelcome, so the Team Syntegrity counterweight cannot be deployed. In those settings the individual instruments still travel — the three-level control-variable check, the pre-control chain as a diagnostic lens, the "who are we not?" question — but the integrated framework, which asks an organization to manage all three of its clocks at once, needs a leadership prepared to look at all three.

Placed beside the others, Schwaninger's is the method that reframes the shared question rather than just answering it. The set asks *who holds the pen, and with what?*:

- **Beer:** the manager, alone, with four hand-drawn charts.
- **Pfiffner:** a confidential core team of three to five, with a variety-measured seven-step process.
- **Hoverstadt:** the consultant, reading an organization by its symptoms.
- **Pérez Ríos:** the analyst, with a checklist, a matrix, a pathology catalogue, and software.
- **Espejo & Reyes:** an enabling viewpoint that convenes stakeholders to name the organization.
- **Espinosa & Walker:** the whole community, building the model together.
- **Schwaninger:** a leadership team steering by three clocks at once, ideally with the many brought in through syntegration.

Schwaninger's distinctive contribution to the question is that *what the pen writes* — the control variables an organization steers by — matters as much as who holds it, and that most organizations are holding the pen over only one of the three pages they are actually being graded on.

What to Borrow

The three-clock check. For any organization, ask what it is actually steering by. If every dial on the dashboard is operative — revenue, cost, cash — the strategic and normative clocks are running unwatched, and the organization is as exposed as the bank in green. Naming the control variables at all three levels is a fast, powerful diagnostic that needs no software.

Pre-control as a lens. When something is failing now, look one logical level up and one time-horizon back: the fast failure is usually the delayed consequence of a slow variable that was run down earlier. "Start earlier, not faster."

Logical level ≠ recursion level. Whenever a debate breaks out over whether strategy or vision "belongs" at the top, reframe it: strategic and

normative management are logical levels to be exercised *at every level*, not floors of a hierarchy. This single distinction ends a whole class of unproductive centralize-the-brain arguments.

"Who are we *not*?" In any identity or ethos discussion, add the negative question. Schwaninger's point is cybernetic, not rhetorical: a sharp preclusion absorbs more variety than a warm affirmation, and tells a team more about how to decide.

What does not transfer for free is the integration. The framework's real power is in binding structure, strategy, finance, ethics and time into one scheme and steering all of them together — and that requires a leadership willing to manage all three clocks, and often the model-based and simulation work the book leans on. Borrow the lenses solo; the full instrument needs a metasystem that wants to see itself whole.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5.

Chapter 7

The Bestiary of Failure

Hoverstadt's Field Guide — Diagnosing Organizations by the Way They Fail

In Brief

Every other method in this set teaches you to build the model. Patrick Hoverstadt's teaches you to recognize the disease. *The Fractal Organization* (2008) is the most practitioner-facing VSM book by a wide margin — written, in its own words, for "practical managers, not academics", kept "as mercifully free as possible from maths, formula and theory" — and its distinctive engine is not a modelling procedure but a **bestiary of twenty-one named structural pathologies**, each described by its symptoms, its underlying structure, and often a real example. The Control Dilemma. Organizational Cancer. The Stray Lamb. Baronies. The Bean Counters. The Death Spiral. A manager who has never heard of Beer can open the book, recognize the exact symptoms afflicting their own organization, and read off the structural cause.

Hoverstadt's premise is that the VSM is a *science* of organization — that "simple laws govern whether an organization will be viable or whether it will die", the way the laws of aerodynamics govern whether a plane will fly — and that most organizational failure is therefore not bad luck or bad people but predictable structural dysfunction. He opens with two client phone calls about restructurings whose failures he could forecast in minutes, and the whole book is an argument that you can learn to see those failures coming.

This chapter describes his method (a light modelling sequence plus pattern-matching plus the pathology catalogue), names what is unique about it, identifies its best application case, and states its limits. It then locates the politics — which Hoverstadt, unusually, addresses head-on: hierarchy as "essentially a religious concept", real decision power as always already distributed regardless of the org chart, and participation defended not as workplace democracy but as a cybernetic necessity.

A note on the title. "Sustainable" here means *viable* — long-lasting, able to survive a changing environment — not ecological. The book is about why organizations die, and how not to.

The Phone Call He Could Answer in Minutes

The book opens with two phone calls. In each, a client describes a restructuring already under way, and in each, Hoverstadt can tell them — quickly, and as it turns out correctly — how it will fail, which functions will break, and which individuals will end up in conflict with which others. Not because he is clairvoyant, but because the proposed structures violate the same laws every time, and violating them produces the same symptoms every time.

His analogy is deliberate and repeated: "we would not dream of flying in a plane… designed by someone who didn't understand the laws of aerodynamics" — yet we routinely restructure organizations by people who have no theory at all of what makes an organization viable, and then treat the resulting failures as surprises (printed pp. 9–10). The org chart, he notes, was "originally invented to attribute blame after a railway accident" — it models the power-and-blame structure and "nothing fundamental" about what the organization actually is or does (printed pp. 42, 65).

The rest of the book is the theory those restructurings lacked, delivered in the form a working manager can use: not "here is the model, now build it", but "here are the ways organizations fail, learn to recognize them."

The Method: Statics, Dynamics, and Pattern-Matching

Hoverstadt distinguishes two uses of the model (Appendix): **analysis** — diagnosing an existing organization, which is "particular to a specific organization… each project is unique", like writing a recipe — and **design**, where building one subsystem (a performance system, a strategic-risk process) affords "much more consistent and prescribed methods", like following a recipe.

Every analysis has two layers. **Statics** are the structures; **dynamics** are the behavioural effects of complexity imbalances between them. The academic instinct is to map the statics first and derive the dynamics; Hoverstadt says real work runs the other way, because "the business problem will be expressed as a dynamic" — something visibly not working — and you trace it back to the structure that produces it. His slogan: "statics drive dynamics and dynamics drive statics." Static analysis is explicitly **pattern-matching** — the VSM used "as a sort of template, against which the real organization… is compared", asking of each thing what its systemic role is, and, crucially, "what it should be doing, but isn't."

He does give a repeatable statics-first sequence (Appendix, steps a–i): **define the identity** of the system under investigation; **unfold the complexity** of its operations; **model the environments** at each level and their complexity drivers; **model the connections** between operations and between environment subsets; **model the coordination** (System 2); **model the management functions** (Systems 3, 4, 5) and their links; **check for missing components and missing links**; **check the complexity imbalances** (the

dynamics); and **work out the implications** — do the diagnosed structural faults match the observed symptoms and point to a fix? The whole is a learning loop: "a model of an organization is a hypothesis that needs to be tested against reality."

He is disciplined about scope. The rule of thumb is to model the system of interest plus one recursion level up and one down. And he names the three scoping mistakes that wreck real projects: a client-imposed scope too narrow to reach the root cause ("both of analysis and politics"); chasing every connection ("no such thing as a complete picture, all models are potentially infinite"); and the "expanding horizon" syndrome, where the scope creeps up recursion levels "until you are faced with having to restructure the whole socio-economic system of the entire country."

The S1 tool. The core design move is the unfolding of operations by **four complexity drivers — technology, geography, customers, and time** — and Hoverstadt's insistence, which he demonstrates with a manufacturing case (work-in-progress down 62%, throughput times down 70%, overdue orders down 82% when cellular replaced functional grouping), that the *order* in which you apply the drivers "can have an absolutely massive impact on how the organization performs." Structure is not neutral plumbing; each structural choice "redraws the boundary between the organization and its environment", opening some strategic options and foreclosing others. His warning is against **"unconscious design"** — the hardware shop that drifted into a supermarket structure whose accidental shape then "dictated the company's strategy."

The subsystem recipes. For design work he offers named, repeatable methods: **Dynamic Performance Management** (measuring not against targets but through the ratios of Actuality, Capability and Potentiality — so that a company beating its targets can still be diagnosed as failing, because the analysts are watching the gap to *potentiality*); **strategic risk** modelled as the breakdown of structural coupling with stakeholders; and **Mosaic Transformation** for change, an eight-step design premised on the fact that different parts of an organization can absorb change at different rates, so the homogeneous "resistance to change" model is simply wrong.

The Twenty-One Pathologies

The book's signature, and its most borrowed contribution, is the catalogue of structural failure modes — "a set of 21 very common systems failures… described by their symptoms, their underlying structure and in some cases by example." A representative sampling conveys the method:

- **The Control Dilemma** — one of the three most common — managers micro-manage under rising complexity, so units must serve both their environment and their managers' demands, and do neither well.

- **Organizational Cancer** — a support function "goes into business on its own" and grows without limit, its normal controls having failed (Hoverstadt's plain-English name for what Beer analyzed as *pathological autopoiesis* and Pérez Ríos catalogues as the autopoietic "beast").
- **The Stray Lamb** — a primary activity that exists but appears nowhere in the management structure, running unmanaged and unseen.
- **Baronies** — sub-units so self-sufficient they refuse the synergy of the whole, hoard resources, and engage in "fratricidal competition."
- **The Bean Counters** — performance management dominates while development is starved; "particularly in Anglo-Saxon cultures", and so pervasive managers can't see it. Its mirror image is **Castles in the Air**, where the future dominates at the expense of delivering today.
- **The Missing Link** — "black holes" in the management structure: whole areas with no performance management or no development, sometimes whole missing levels.
- **Silo Decisions** — strategy made without cross-discipline conversation (his examples include the Sinclair C5 and UK carmakers ceding the 4×4 market).
- **The Death Spiral** — a "meta-archetype" that chains several others together: governance fails to balance strategic decisions → environmental change goes unnoticed → operational response turns erratic → strategic collapse → bunker mentality → "organizations in this state only survive by luck."

The three he names as most common — missing coordination links, the Control Dilemma, and missing primary-activity representation — are the ones a diagnostician should test for first. The catalogue works because it is symptom-first and memorable: managers recognize their own organization in the names before they have understood a line of cybernetics, and the recognition is the hook that carries them to the structural cause.

The catalogue is original; its content is not, and Hoverstadt would not claim otherwise. Every entry names a recurring violation of laws Beer stated, and the origin should be said plainly: the founding pathology of the whole genre is Beer's own. *The Heart of Enterprise* diagnoses **pathological autopoiesis** — "Enterprises need to be autopoietic; their pathology is to turn self-production into an end-in-itself" (ch. 15, "Life and Death", p. 410) — and Beer even supplies the cancer image himself, offering "a cancerous tumour" as a biological instance of the disease (p. 412). Organizational Cancer, the bestiary's most memorable name, is thus a development of Beer's own imagery; the Control Dilemma restates a variety imbalance on the command axis; the Missing Link, a violated First Axiom. What the catalogue adds is a way in: a route to the fault that does not pass through the theory.

What Is Actually Unique Here

Diagnosis by symptom. No other VSM method organizes itself around a named catalogue of failure modes. It inverts the usual order — instead of "learn the model, then apply it", it is "recognize the disease, then learn the structure behind it" — which is why it reaches practising managers that the academic treatments do not. The design principle deserves to be stated explicitly: the catalogue puts the pathology directly into the manager's hand, with no mandatory understanding of the method. Beer's diagnosis is generative and therefore gated — nothing becomes visible until the model has been built and the laws learned. Hoverstadt removes the gate. Recognition comes first, and the cybernetics behind the recognition is the payoff, not the prerequisite: you do not need to know what a variety imbalance is to know that your support function has gone into business on its own and will not stop growing. The price is paid in judgment rather than access — without the model, a reader can recognize a pathology but cannot reliably tell which one is really present, or a cause from a consequence — which is why the eye, not the list, is what the book is finally trying to teach.

Plain English, and the anti-fuzzy stance. Hoverstadt strips out the mathematics, the neurophysiology and most of the jargon (relegated to a glossary), while insisting that in his diagrams "the lines all do mean something" — a pointed rejection of the vague boxes-and-arrows that pass for systems thinking elsewhere. It is the most accessible serious VSM book in print.

"Fractal" instead of "recursion", and structure-drives-behaviour. He generally prefers *fractal* to Beer's *recursion* — the same design rules replicated at every scale, so that "big systems have the same systemic properties as small ones; they work and fail in exactly the same ways." The payoff is his central claim that structure causes behaviour predictably: get the structure wrong and you can forecast the dysfunction, which is what makes the bestiary possible at all.

The critique of conventional strategy. His demolition of linear strategy (the "vision → mission → strategy → targets → measures" cascade, of which "over 90% of strategic plans are never implemented") and his reframing of strategy as an *emergent property* built from "conversational processes between different levels" is the most usable strategic-management application of the VSM available.

Strengths, Best Case, Limits

Strengths. It is the most immediately usable VSM book for a working manager or consultant — accessible, symptom-first, and grounded in a rich bank of real (anonymized) consulting cases. It is strong precisely where the academic treatments are thin: on the *dynamics* of structure, on why bad structure produces predictable failure, on the resource-performance

bargain, and on strategy, strategic risk and governance as concrete VSM applications.

Best application case. The book is at its best in the hands of a practising manager or internal consultant who is *already seeing symptoms* — recurring turf wars, a restructuring that isn't taking, performance games, a strategy that never lands — and wants a fast, jargon-light structural diagnosis they can act on. It is the ideal entry point to the whole VSM tradition, and the right first tool when the presenting problem is a live dysfunction rather than a clean-sheet design. It is also unusually good for the manager who suspects, correctly, that their org chart is lying to them.

Limits. Three. It is deliberately light on the formal cybernetics — transduction, the algedonic channel as a named mechanism, the mathematics of variety — so a reader who needs rigor or a defensible formal model will need Beer, Espejo or Pérez Ríos alongside it. There is no software or standard worksheet (contrast Pérez Ríos's VSMod® or Espejo's recursion/functions table); the method is a modelling sequence plus a diagnostician's judgment. And the cases are anonymized and the failure statistics ("90%", "70%") cited loosely — persuasive in the room, unfalsifiable on the page. It is a field guide, not a proof.

Part I — The Political Environment the Method Was Built For

Hoverstadt is the most direct of the VSM writers about organizational power, and his position is bracing: real decision power is *already distributed*, regardless of what the org chart claims, and the only question is whether you design to use that fact or waste energy denying it.

Hierarchy, he writes, is "essentially a religious concept" — the people at the top imagined as closer to God — which reached its apogee in Taylorism and survives mainly because it flatters top decision-makers and because dismantling it "requires the very people who benefit to agree" (printed pp. 64–65). It is also, he argues, "intrinsically unstable" and "often an illusion." His anchoring case is a service organization whose vaunted "command and control" turned out on inspection to be "command and *ignore*": the middle managers had quietly "wrested control of their own processes", exercising an autonomy that "was not gifted… not 'empowered'… They just did it." His conclusion is flat: "People have autonomy to act whether we like it or not" (printed pp. 66–67). Even prisons, he notes, "function on a consensual basis."

This gives his politics a distinctive shape. Participation, in Hoverstadt, is a **cybernetic necessity, not an ideology**: the resource-performance bargain and strategy-making must be "communicative and participative" because no small management group can possibly have requisite variety against the complexity it faces, and because involvement is inversely correlated with resistance. He grounds this in Chapter 2's treatment of self-fulfilling

prophecies — McGregor's Theory X and Y, the zero-sum "nice guys finish last" belief — arguing that management models are self-confirming, so a theory that treats people as needing control produces organizations that require it. But — and this distinguishes him sharply from Espinosa and Walker, or from Schwaninger's normative apparatus — he does *not* frame this as workplace democracy or as a matter of rights. Autonomy and distributed decision-making are justified on **variety grounds**, full stop. The argument is engineering, not ethics.

The application domain is deliberately universal — "any organization of any size in any sector... private or public sector, ruthlessly self-serving or wholly altruistic" — though the cases skew private (manufacturing, IT, engineering, retail, finance). The public sector is nonetheless prominent: the opening crisis is a public body under a regulator, and recurring cases include the health service, a government roads agency, and the nuclear-waste decision system spanning some twenty organizations. Supply chains and virtual joint ventures appear; governance cases include Enron and Barings.

The best political environment for this method, then, is almost any organization willing to accept an uncomfortable premise: that its formal authority structure is not where its real control lives, and that good structure means designing *with* the distributed decision-making that already exists rather than against it. Hoverstadt does not require a commissioner to hold progressive values — his argument works for the "ruthlessly self-serving" organization exactly as well as the altruistic one, because it is about variety, not virtue. What it requires is a willingness to stop pretending the org chart is the organization. Where that willingness exists, the diagnosis lands. Where a leadership is committed to the fiction of top-down control, the method will keep telling them the control they think they have is "command and ignore."

Part II — The Internal Power Constellation It Presupposes

Hoverstadt's method is, in execution, the **consultant's** — a skilled diagnostician reading an organization by its symptoms — but its entire content is a claim about power that undercuts the usual consultant's assumption of a single decision-maker who can implement the findings.

The pen is the diagnostician's; the authority is diffuse. The modelling and pattern-matching are done by an analyst (internal or external) with the eye to recognize the pathologies. But the diagnosis itself keeps returning the verdict that *no one is actually in the control position the chart assigns them* — that autonomy is exercised at every level whether sanctioned or not. So the method presupposes not a powerful sponsor who can push a redesign through, but an organization in which the redesign must work *with* the real, distributed loci of decision, because the formal ones are partly illusory.

Scope is where power bites. Hoverstadt is explicit that the most damaging constraint on a diagnosis is a client-imposed scope too narrow to reach the root cause — a limitation he calls "both of analysis and politics." The sponsor who commissions a diagnosis but fences off the area where the real dysfunction sits is exercising precisely the kind of self-protective power the method is built to expose. The internal precondition for the method to work, therefore, is a sponsor secure enough to let the analysis follow the symptoms wherever they lead — including into their own function.

Trust, built by monitoring, is the load-bearing relationship. His treatment of the resource-performance bargain and of System 3* monitoring turns on trust: more autonomy plus reliable reporting produces better delivery, but only if monitoring is sporadic, in-depth and honest enough that both sides believe the reports. The method presupposes a management willing to build that trust rather than substitute surveillance for it — the same choice, in different words, that every method in this set eventually forces.

Where the internal preconditions fail. They fail where the sponsor fences the scope to protect themselves; where leadership is too committed to the authority-and-obedience picture to accept that control is already distributed; and where monitoring is used to control rather than to build trust, so the performance bargain collapses into gaming. In those settings the pathology catalogue still *describes* the organization accurately — indeed it will name exactly what is wrong — but the organization cannot act on the description, because acting would mean surrendering the fiction the dysfunction protects.

Placed beside the others, Hoverstadt is the diagnostician of the set, and his distinctive answer to the shared question is that the question partly answers itself:

- **Beer:** the manager, alone, with four hand-drawn charts.
- **Pfiffner:** a confidential core team of three to five, with a variety-measured seven-step process.
- **Pérez Ríos:** the analyst, with a checklist, a matrix, a pathology catalogue, and software.
- **Espejo & Reyes:** an enabling viewpoint that convenes stakeholders to name the organization.
- **Schwaninger:** a leadership team steering by three clocks at once.
- **Espinosa & Walker:** the whole community, building the model together.
- **Hoverstadt:** the consultant, reading an organization by its symptoms — while insisting that whoever *holds* the pen, the *power the pen describes* is already spread through the whole organization, and always was.

That is Hoverstadt's contribution to the question. The others ask who should hold authority over the model. Hoverstadt observes that authority over the *organization* is never where the chart says it is — which is the first

thing a good diagnosis reveals, and the last thing a threatened leadership wants to hear.

What to Borrow

The pathology catalogue as a first pass. Before building any model, walk a struggling organization against the named failure modes — Control Dilemma, Organizational Cancer, Stray Lamb, Baronies, Bean Counters, Missing Link, Silo Decisions, Death Spiral. It is the fastest structured diagnosis available, and its symptom-first framing gets managers nodding before they have learned any theory.

Start from the dynamic, work back to the structure. The presenting problem is always a behaviour — something not working. Trace it to the structural imbalance that produces it, rather than starting from a blank org model. "Statics drive dynamics and dynamics drive statics."

Mind the order of the complexity drivers. When unfolding operations, the sequence in which you apply technology, geography, customers and time is not cosmetic — it changes what the organization can do. Choose it consciously, and beware the "unconscious design" whose accidental structure quietly dictates the strategy.

Measure potentiality, not just targets. Dynamic Performance Management's three ratios — actuality, capability, potentiality — catch the organization that is beating its targets while falling behind its environment, which no target-based dashboard can see.

Say the quiet part about the org chart. In any structural conversation, put Hoverstadt's observation on the table: the chart models power and blame, not what the organization does, and the real control is already distributed. It reframes the whole discussion from "who should be in charge" to "how does this organization actually work" — which is the only question a diagnosis can answer.

What does not transfer for free is the eye. The bestiary is memorable, but recognizing which pathology is really present — and which symptom is a cause versus a consequence — is a diagnostician's skill built on many cases. Borrow the catalogue as a checklist; the judgment behind it is the thing the book is actually trying to teach, and it takes longer than one read.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. (Ch. 15, "Life and Death", pp. 403–414: pathological autopoiesis.)
- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Pérez Ríos, J. (2010). "Models of organizational cybernetics for diagnosis and design." *Kybernetes*, 39(9/10), 1529–1550.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5.

Part II

Part II — Choosing a Method

From Seven Methods to One Decision

Having seven methods is not the same as knowing which to use. Part I profiled the shelf; Part II is for the moment you stand in front of it with a real organization and a real problem.

The capstone that follows is practical. It lays the seven side by side across the axes that separate them — who holds the pen, the best application case, effort, participation, tooling, what "sustainable" is taken to mean, the political stance each carries — and turns the comparison into a decision guide keyed to your first honest question about the organization. Its claim is that choosing well is mostly a matter of matching the method's placement of authority to how power and agreement are really distributed where you work.

One thing the capstone does not do is close the argument. Every method in Part I answers a political question in structure — where authority over the model sits — while mostly declining to name it as one. Naming it is the job of the coda, *Designing Freedom, Again*, which closes the book after the showcases of Part III: Beer's own late argument about freedom, the way AI moves organizational structure through System 2, and the distribution of human agency as the decisive variable. Read the capstone now; save the coda for last.

- **Choosing a Method** (ch. 8): the capstone comparison and decision guide across all seven — the tables, the "if your problem is X, start with Y," and how to combine methods.

If Part I is the shelf and Part III is the machinery running, Part II is where you decide which instrument to pick up — and remember why the choice was never merely technical.

The capstone draws on all seven method chapters of Part I. The coda that closes the book draws additionally on Beer's Designing Freedom *(1974) and the working paper* Human Agency in the Age of Artificial Intelligence *(2026); Showcase 6 in Part III is its practical companion.*

Chapter 8

Which Pen, and When

Choosing Among the Seven VSM Methods — a Comparison and a Decision Guide

In Brief

Seven books, one model. Each of the preceding chapters takes Stafford Beer's Viable System Model and profiles a distinct method for putting it to work — Beer's own do-it-yourself handbook, Pfiffner's seven-step design process, Espinosa and Walker's self-transformation methodology, Pérez Ríos's software-carried instrument, Espejo and Reyes's identity-first VIPLAN, Schwaninger's two-axis integration, and Hoverstadt's symptom-first field guide. A reader who has worked through all seven arrives at a practical question the individual chapters answer only for their own subject: *given a real organization and a real problem, which one do I actually reach for?*

This chapter answers that question. It does not re-profile the methods; it lays them side by side across the axes that discriminate between them — who holds the pen, the best application case, the unique contribution, effort and scale, participation, tooling, what "sustainable" is taken to mean, whether the method leads with diagnosis or design, and the political stance each carries — and then turns the comparison into a decision guide keyed to the problem in front of you.

One finding organizes everything that follows, and it is the thread the whole set has been pulling: **these seven are not rival readings of the VSM. They are different placements of authority over the model.** The model is the same in all of them. What changes is *who holds the pen* — one manager, a confidential core team, an enabling facilitator, an expert analyst, a whole community, a top team steering by three clocks, a consultant reading symptoms — and choosing well is mostly a matter of matching that placement to how power and agreement are actually distributed in your organization, and matching the method's leading move to the problem you actually have. Get those two matches right and any of the seven will repay the effort. Get them wrong and the most rigorous method in the world produces, in Espinosa and Walker's phrase, a very well-attended failure.

The Comparison

Two tables read the seven against nine axes — the first holds the political axes, the second the practical ones. It is a map, not a verdict: the "best case" column is where each method is *strongest*, not the only place it works, and several of the methods overlap productively (the next section takes up combining them). Page-level grounding for every cell is in the chapter cited in the first column.

Table 1 — The pen and the politics.

Method (chapter)	Who holds the pen	Participation	Political stance	"Sustainable / viable" means
Beer, *Diagnosing the System* (ch. 1)	One manager, alone or with an aide	Essentially none, by design	Technocratic: autonomy as a matter of variety "rather than political theory"; power named, then bracketed	*Viable* = capable of independent survival (not ecological)
Pfiffner, *The Neurology of Business* (ch. 2)	A confidential core team of three to five; decision reserved to leadership	Instrumental: Syntegration in Phase B, but the board reserves the decision	Technocratic ("wars of ideas... unnecessary"); implicit Swiss federalism; co-determination left as a gap	*Viable* = mastery of complexity
Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3)	The whole community, co-constructing the model	Maximal and constitutive — and only counts if it carries real decision rights	Explicitly progressive: commons governance, anti-neoliberal critique; needs a non-repressive setting	*Ecological-societal sustainability* — the one method where the word means what it says

Method (chapter)	Who holds the pen	Participation	Political stance	"Sustainable / viable" means
Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4)	The expert analyst, at a desk, with software	Quarantined: Team Syntegrity is a separate, optional chapter	Technocratic; public-sector home; no place to represent a contest of interests	*Viable* = long-term structural survival (explicitly *not* ecological)
Espejo & Reyes, *Organizational Systems* / VIPLAN (ch. 5)	The *enabling viewpoint*: an expert catalyzes; participants settle the naming	Genuinely participative and constitutive (naming and discretion workshops)	Worn on the surface: an ethics of autonomy (von Foerster); hierarchy as "bad cybernetics"; democratic legitimacy in the public cases	*Viable*, with identity foregrounded
Schwaninger, *Intelligent Organizations* (ch. 6)	The top team and its advisers (expert-led in execution; participative in ideal)	Taken seriously and critically; Team Syntegrity as "an architecture for democracy" — necessary, not sufficient	Built in, not bracketed: critique of short-termism and financialization; Mondragón; the triple bottom line	*Viability beyond survival* — development and legitimacy, not just endurance

Method (chapter)	Who holds the pen	Participation	Political stance	"Sustainable / viable" means
Hoverstadt, *The Fractal Organization* (ch. 7)	The consultant/diagnostician — while insisting real power is *already* distributed	A cybernetic necessity on variety grounds, not an ideology	Direct: hierarchy as "a religious concept" and often an illusion; works for the "ruthlessly self-serving" and the altruistic alike — engineering, not ethics	*Viable / long-lasting* (not ecological)

Table 2 — The practicalities.

Method (chapter)	Best application case	Unique contribution	Effort & scale	Tooling	Leads with
Beer, *Diagnosing the System* (ch. 1)	A leader who wants to see an enterprise's deep structure themselves — fast, private, unaided — and large, bewildering systems where recursion collapses the detail	The source method; genuinely do-it-yourself; the parsimony of recursion; four fill-in Charts	Minutes to weeks; solo; hand-drawn on a wall	None (photocopier, scissors, paste)	Diagnosis

Method (chapter)	Best application case	Unique contribution	Effort & scale	Tooling	Leads with
Pfiffner, *The Neurology of Business* (ch. 2)	A mid-to-large firm under one accountable leadership facing a real structural question (segmentation, integration, span of control), able to decide and implement	Settles autonomy questions by *measuring variety* instead of arguing; the "third dimension"; a seven-step manual	Elastic — "5 min… to several months"; team of two to five; ~9-month reorg	Tools, tests, worksheets; no software essential	Design
Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3)	Co-ops, eco-communities, participatory programmes, commons — where those inside own the structure and there is a real forum to decide	The VSM as a *meta-language* the organization learns and keeps; the consultant refuses to diagnose; viability extended to the ecological niche	Months of workshop cycles; whole membership	Self-chosen indicators; no central software	Diagnosis (self-produced)

Method (chapter)	Best application case	Unique contribution	Effort & scale	Tooling	Leads with
Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4)	Large, multi-scale bodies with no single natural hierarchy — universities, regions, health systems — where sustained expert work and documentation are wanted	Built *around an instrument* (VSMod®); multi-criteria recursion as a sphere; a named pathology catalogue; the exists/quality/performs triple test	Heavy; sustained expert work over time	**VSMod® — central** (the only method built around a program)	Both, symmetrically
Espejo & Reyes, *Organizational Systems* / VIPLAN (ch. 5)	An organization whose identity is contested or unclear — a merger, a regulator, a public body pulled between mandates	Identity-first: naming the organization *is* the first and most consequential act; TASCOI; a second-order-cybernetic epistemology	Facilitated workshops over time	VIPLAN software (the build from the recursion/functions table is handed off to it)	Diagnosis (identity-first)

Method (chapter)	Best application case	Unique contribution	Effort & scale	Tooling	Leads with
Schwaninger, *Intelligent Organizations* (ch. 6)	A top team at an inflection point — "winning on the metrics and losing somehow anyway" — with eroding value potentials or legitimacy	A *second axis*: three logical levels (operative/strategic/normative) orthogonal to recursion; pre-control; viability *beyond* survival; five models integrated	Not a quick fix; abstract; model-based	System Dynamics simulation (one of the five models)	Both (diagnosis of what you steer by)

Method (chapter)	Best application case	Unique contribution	Effort & scale	Tooling	Leads with
Hoverstadt, *The Fractal Organization* (ch. 7)	A practising manager or internal consultant *already seeing symptoms* — turf wars, a reorg not taking, performance games — who wants a fast, jargon-light read	Diagnosis by symptom: a bestiary of twenty-one named pathologies; plain English; statics/dynamics	Fast, light, judgment-based; system plus one level up and down	None (no software, no standard worksheet)	Diagnosis

Three readings *down* the columns discriminate more than any single row.

The pen-spectrum is the spine. Run the "who holds the pen" column from the leanest constellation to the widest and the whole set arranges itself: Beer's *one manager* → Pfiffner's *confidential core team of three to five* → Hoverstadt's and Pérez Ríos's *single expert diagnostician* (the consultant reading symptoms; the analyst driving software) → Espejo and Reyes's *enabling viewpoint* that shares the pen with convened stakeholders → Schwaninger's *top team steering by three clocks* → Espinosa and Walker's *whole community* holding the pen together. This is not a quality ranking. It is a ranking of *how many people must agree* before the method's output can become the organization's reality — and that, more than rigor or elegance, is what determines whether a method will take.

"Sustainable" means two different things, and the difference is a trap. Four of these books carry "sustainable" or "sustainability" in or near their titles, and only *one* — Espinosa and Walker — uses it to mean ecology, climate and the biosphere. For Pérez Ríos, Hoverstadt and the rest, "sustainable" is a synonym for *viable*: a structural property of an organization that keeps its identity through change. A reader who picks up Pérez Ríos's *Design and*

Diagnosis for Sustainable Organizations expecting an environmental method has misread the genre; a reader who picks up Espinosa and Walker expecting only structural viability has missed the point. When the word appears, check which sense is meant before you choose.

Effort and tooling move together, and neither tracks rigor. The lightest methods — Beer's charts, Hoverstadt's bestiary — carry no software and can produce a usable read in an afternoon. The heaviest — Pérez Ríos's instrument-built analysis, Schwaninger's five-model integration — reward patience over dispatch and lean on software or simulation. But light does not mean shallow: Hoverstadt's symptom-first read is often the fastest route to the true structural cause, and Beer's is the primary source everything else interprets. Effort should be chosen to fit the decision at stake, not treated as a proxy for seriousness.

The Decision Guide

The reliable way in is not "which method is best" but "what is the *first honest question* about this organization." Each method leads with a different first question, and matching yours to theirs is most of the choice.

Start from the live problem.

If you want to *understand the deep structure yourself* — quickly, privately, without hiring anyone — start with **Beer's four Charts** (ch. 1). It is the cheapest to begin and the least mediated, and working it once by hand teaches the distinctions every other method assumes.

If something is *visibly broken* — recurring turf wars, a restructuring that will not take, performance games, a strategy that never lands — start with **Hoverstadt's pathology catalogue** (ch. 7). It is symptom-first, so it gets managers nodding before they have learned any theory, and it is the natural entry point to the whole tradition when the presenting problem is a live dysfunction rather than a clean-sheet design.

If you are *restructuring a mid-to-large firm* under a single accountable leadership that can decide and implement, start with **Pfiffner's seven steps** (ch. 2). It is built for exactly that commissioner, it measures where autonomy should sit instead of arguing about it, and its nine-month, confidential-core-team design is a coherent plan for getting from diagnosis to a changed org chart.

If you *cannot agree what the organization even is*, or for whom — a merger, a regulator, a public body pulled between mandates — start with **Espejo and Reyes's VIPLAN** (ch. 5). Its whole premise is that naming the organization is the first and most consequential act, and its workshops are built to settle a contested identity before anything downstream is drawn.

If the *people inside must own the change* — a co-operative, an eco-community, a commons, any setting where legitimacy to restructure is distributed rather than held — start with **Espinosa and Walker's self-transformation method** (ch. 3). It is the only method that makes the

organization the author of its own diagnosis, and the only one whose "sustainability" reaches the ecological niche.

If the organization is a *large, multi-scale body with no single natural hierarchy* — a university across its campuses and region, a health system, a territorial administration — and you want rigor and documentation over speed, start with **Pérez Ríos's Viable System Method** (ch. 4). Its multi-criteria recursion and its software are built precisely for the system that can be legitimately cut into levels more than one way.

If you are *winning on the metrics and losing somehow anyway* — the operative numbers look fine but value potentials or legitimacy are quietly eroding — start with **Schwaninger's three clocks** (ch. 6). Its second axis asks what you are actually steering by, and it is the method that catches the organization graded on three pages while writing on only one.

A decision tree. The same guidance, as a quick reference. Take the first branch that fits; the branches are ordered so the more specific conditions come first.

```
What is the first honest question about this organization?

  "We don't agree what this organization even is, or for whom."
     → Espejo & Reyes - VIPLAN (ch. 5)   [identity-first]

  "Something is visibly broken and I want a fast read."
     → Hoverstadt - the pathology bestiary (ch. 7)   [symptom-
first; also the best entry point]

  "We're winning on the metrics and losing somehow anyway."
     → Schwaninger - the three clocks (ch. 6)   [what are we steering by?]

  "The people inside must own the change."
     → Espinosa & Walker - self-transformation (ch. 3)   [only method where
           …and sustainability means ecology, not just survival? → confirms

  "It's a large, multi-scale body with no single natural hierarchy,
    and I want rigour + documentation."
     → Pérez Ríos - the Viable System Method + VSMod® (ch. 4)   [multi-
criteria recursion]

  "We're restructuring a mid/large firm under one leadership
    that can decide and implement."
     → Pfiffner - the seven steps (ch. 2)   [measures variety; ~9-
month reorg]

  "I want to understand the deep structure myself, fast and private."
     → Beer - the four Charts (ch. 1)   [the source; do-it-
yourself]
```

Two cross-checks refine any branch. First, **who must agree?** If the honest answer is "one accountable leader," the lean methods (Beer, Pfiffner, Hoverstadt) fit; if it is "a distributed set of stakeholders whose consent is required," you need the participative ones (Espejo and Reyes, Espinosa and Walker) or you will produce a diagnosis nobody adopts. Second, **is the identity settled?** If not, no amount of structural rigor will help until it is named — which is why VIPLAN sits at the top of the tree.

Combining and Sequencing

The methods are not mutually exclusive, and the strongest engagements borrow across them.

The lightest two make natural *first passes*. Walk a struggling organization against **Hoverstadt's pathologies**, or work **Beer's four Charts** by hand, to get a fast structural read before committing to a heavier method. Either will surface the questions the deeper work then answers.

Two of **Pérez Ríos's** instruments are portable to *any* method and need no software: the **exists / quality / performs triple test**, applied to every component, converts a diagram everyone nods at into findings some things fail; and the **pathology catalogue** doubles as a checklist you can run on paper. Borrow both regardless of which method holds the frame.

The two pathology catalogues: one origin, opposite constructions. Since the set contains two catalogues — Hoverstadt's bestiary and Pérez Ríos's coded taxonomy — it is worth being precise about their relationship, because it is closer than it looks. Both descend from Beer. The founding pathology of the genre is his: *The Heart of Enterprise* diagnoses **pathological autopoiesis** — "Enterprises need to be autopoietic; their pathology is to turn self-production into an end-in-itself" (ch. 15, p. 410) — and Beer himself offers "a cancerous tumour" as its biological instance (p. 412), so that Hoverstadt's Organizational Cancer develops Beer's own imagery and Pérez Ríos's autopoietic "beasts" (his pathologies II14 and II16) give the same disease its taxonomic entry. The lineage is documented, not inferred: Pérez Ríos cites Beer's own pathology paper — "The viable system model: its provenance, development, methodology and pathology" (1984/1989) — as the analytical source of his catalogue. Beer compiled no catalogue himself because his diagnosis is *generative*: build the model, test it against the Principles and Axioms, and derive the fault — including one that has never occurred before and has no name. The price of that generality is that nothing is visible until the model has been learned, and no shared vocabulary results. The catalogues are what happened when practitioners noticed that the violations recur; what each adds is not content but *access* — a way of finding the fault without deriving it.

The two compilations then part ways on construction, and the difference tells you which to reach for. **Hoverstadt's is symptom-first**: named by how failure looks and feels, in plain English, so that the pathology sits directly in the *manager's* hand with no mandatory understanding of the

method — recognition first, theory as the payoff — and with the Death Spiral it even models how pathologies chain into a failure trajectory over time. What it cannot transfer is the judgment: without the model, a reader can recognize a disease but not reliably tell which one is present, or a cause from a consequence. **Pérez Ríos's is component-first**: organized by the model element that fails (structural, functional system by system, informational), coded, and exhaustive by construction, so an *analyst* cannot skip a component — but unusable without the model in hand, and static, with no account of how faults compound. Coverage gaps run both ways: Pérez Ríos treats the information and communication pathologies systematically where Hoverstadt leaves the algedonic channel implicit; Hoverstadt covers the strategy and governance dynamics a component walk-through underweights. And **Espinosa and Walker's lack of a catalogue is a refusal, not a gap**: a pre-named catalogue is a diagnosis the consultant carries into the room, and their method's defining commitment is that the naming of the disease stays with the community that has it. A pathology catalogue is itself a pen-holding device — from Beer's laws, to Pérez Ríos's analyst's checklist, to Hoverstadt's bestiary, the history of the genre is the history of the pen moving outward, from the model's author to anyone who can recognize their own organization in a well-chosen name.

Team Syntegrity — Beer's non-hierarchical group format — recurs across the set as the participation instrument several methods bolt on: Pfiffner uses it for Phase B, Pérez Ríos devotes a chapter to it, Schwaninger calls it "an architecture for democracy." If your chosen method is expert-led but the decision needs legitimacy, a Syntegration is the standard way to add participation without redesigning the whole approach.

And **Schwaninger's** framework is explicitly an *integration* — it already sets the VSM alongside System Dynamics and normative-management theory — so it is the natural home when you want the structural method to sit inside a larger steering framework rather than stand alone.

One caveat governs all of this, and every chapter in the set arrives at some version of it: you can borrow the *tools* solo, but the participative methods' real yield needs the room. TASCOI run as a desk exercise gives you a naming; run as a convened workshop it gives you an identity the organization actually shares. The self-transformation stages studied alone teach the model; run with the community they change the organization. Borrow the instruments freely — but do not mistake the instrument for the method that gives it its power.

What to Borrow

Name who holds the pen before you name the method. The single most consequential decision is not which of the seven you use but where you

place authority over the model — one manager, a core team, an enabling facilitator, an expert analyst, the whole community. Decide that consciously, against how agreement is actually distributed in your organization, and the method almost picks itself.

Match the method's leading move to your first honest question. Identity contested → VIPLAN. Something broken → Hoverstadt. Steering unclear → Schwaninger. Ownership must be internal → Espinosa and Walker. No natural hierarchy → Pérez Ríos. Firm to restructure under one leadership → Pfiffner. Structure to see for yourself → Beer. The methods differ less in what they can analyze than in what they lead with.

Read "sustainable" twice. When the word appears in a method's title, establish whether it means *ecological sustainability* (Espinosa and Walker) or *structural viability* (everyone else) before you build an engagement on it. The two are not the same programme, and the mismatch is the most common misreading in the literature.

Start light, then go heavy only where the decision warrants it. Beer's charts or Hoverstadt's bestiary will give you a usable read in an afternoon; reserve Pérez Ríos's instrument-built analysis or Schwaninger's five-model integration for decisions whose stakes justify the effort. Effort is a cost to be matched to the question, not a proxy for rigor.

The failure mode is choosing without noticing you are choosing. Every method encodes a politics — most bracket the question of whose interests the structure serves; a few put it on the surface. The danger, as the set has argued throughout, is not picking the "wrong" politics but adopting one unawares because the method's framing made the choice feel technical. Choose the pen-holder, the leading question, and the stance on purpose. That is the whole discipline this set has been teaching, gathered into a single decision.

The Method Set

This chapter is the capstone of the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I.

References

- Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. (Ch. 15, "Life and Death", pp. 403–414: pathological autopoiesis.)
- Beer, S. (1984). "The viable system model: its provenance, development, methodology and pathology." *Journal of the Operational Research Society*, 35(1), 7–25. Reprinted in Espejo, R. & Harnden, R.

(Eds.) (1989), *The Viable System Model: Interpretations and Applications of Stafford Beer's VSM*. Chichester: Wiley.

- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Pfiffner, M. (2022). *The Neurology of Business: How the Viable System Model Helps Leaders to Manage Complexity*. Cham: Springer (Management for Professionals). ISBN 978-3-031-14259-8.
- Pérez Ríos, J. (2010). "Models of organizational cybernetics for diagnosis and design." *Kybernetes*, 39(9/10), 1529–1550. (Pathologies II14/II16: autopoietic "beasts".)
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.
- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Schwaninger, M. (2009). *Intelligent Organizations: Powerful Models for Systemic Management* (2nd ed.). Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-540-85161-5.

Chapter 9

Coda: Designing Freedom, Again

The Viable System, Artificial Intelligence, and the Agency Divide

In Brief

The seven methods in this set each answer the question that organizes it — *who holds the pen?* — and each, in doing so, answers a political question in structure while, almost to a book, declining to name it as one. I have called that the technocratic move, and noted each time that it is a move rather than a fact. What I have not done until now is name the one book in the tradition where the politics is said out loud.

That book is Stafford Beer's *Designing Freedom*. It is the shortest and most political thing he wrote — five broadcast lectures rather than a treatise — and its argument is the missing half of everything the method chapters describe. Freedom, Beer says there, is not a property of technology and not a vague ideal. It is a property of organizational design, and therefore a thing that is either built in or engineered out by whoever holds the pen. The same cybernetic tools can be assembled into what he calls a *liberty machine* or into an instrument of oppression, and nothing in the tools themselves decides which.

Artificial intelligence is what makes that fifty-year-old argument suddenly urgent. As the introduction to this set argued, the identical capability that can flatten an organization can also recentralize and surveil it — and structure alone cannot tell you which future you are building. This chapter supplies the criterion that structure cannot. The deciding variable, I will argue, is not the autonomy of the machine but the distribution of *human agency* — and there is now a way to see that distribution (a five-dimensional model of agency) and a way to check it before you deploy (an agency impact assessment). This is the thesis the whole set has been circling, and this coda says it in full.

The Book Where Beer Said the Quiet Part

Designing Freedom began as the 1973 Massey Lectures, delivered for Canadian radio, and it reads nothing like the trilogy the first chapter of this set profiles. There are no charts, almost no jargon, and no worksheet. There is instead a sustained public argument, made by the inventor of the Viable

System Model, about why our institutions fail and what freedom would have to mean if we wanted to design for it.

Beer's premise is that contemporary institutions have become machines for producing results that nobody actually wants — outcomes no participant chose, generated reliably by the structure itself. His diagnosis is cybernetic. Institutions survive by attenuating the variety of the people they serve: they reduce the citizen, the patient, the worker to the handful of states the system can process, and they call the residue noise. Freedom, in this account, is not the absence of constraint. It is the *residual variety* left to a person or a unit to act on their own account — the room to do otherwise. That reframing matters, because it makes freedom an operational quantity, something a design either preserves or destroys, rather than a slogan to be affirmed and ignored.

From there comes the claim the method chapters keep declining to make. The tools of management cybernetics — and in 1973 Beer means the computer above all — are radically double-edged. Assembled one way, they build what he names a *liberty machine*: a system that pushes decision down to where the variety actually is, that lets subsystems act, that uses its informational power to enlarge the room people have to act. Assembled the other way, the very same tools build an apparatus of centralized control, in which the computer becomes the most efficient instrument of oppression ever devised. Beer is emphatic that the technology does not choose. Someone chooses. Freedom has to be *designed* — deliberately, against the institutional default, by whoever is holding the pen.

This is why *Designing Freedom* belongs at the end of this set rather than outside it. Throughout the seven chapters I have treated the political question — whose interests does this structure serve? — as something the methods bracket and I have to supply from outside. That framing was a convenience, and it was slightly wrong. The question is not outside the tradition. It is native to it, asked and answered by the model's own author, in the one book where he stopped describing viable structure and started asking what it is *for*. Bringing *Designing Freedom* back in is not importing politics into the VSM. It is restoring the VSM's own.

What Beer could not have specified, in 1973, was the machine that would make his warning concrete. He was arguing about mainframe computing and centralized databases. The tool that now sits exactly on the fault line he described — liberty machine or instrument of oppression, and the design decides — is artificial intelligence.

The Mechanism: How AI Moves the Structure

To see why AI lands on Beer's fault line, start with what hierarchy is *for*, in variety terms. By Ashby's Law a regulator must command at least as much variety as the thing it regulates, and a human node has strictly bounded variety — limited attention, limited channels, limited time. Full lateral coordination does not scale: a network of n units, each talking to every

other, has on the order of n^2 connections, and no human node can hold that many live relationships at once. Hierarchy is the classical workaround. It attenuates variety by aggregating and summarizing on the way up, and it cuts the connection count from roughly n^2 toward n by routing communication through a command axis. Beer put the human limit plainly in the source handbook: senior management is "in principle unable to entertain the variety generated by any one of its subsidiary viable systems," and the beginnings of decentralization lie in that fact "rather than in political theory." Strong hierarchy, in other words, is partly a symptom — of scarce variety-handling capacity at the top and expensive coordination across the middle.

Artificial intelligence relaxes exactly those scarcities, and it does so in three distinct ways. It amplifies the variety each node can handle, so a single unit can sense, interpret and respond to more complexity locally than a person could alone. It lowers the cost of lateral coordination directly, because a shared model can keep peers aware of each other's relevant state without a manager aggregating and relaying it. And it removes the need to attenuate by summary at all: instead of sending pre-digested reports upward, any node can query raw, high-resolution state on demand. Each of those undercuts a distinct reason for building vertical structure in the first place.

Put it in Beer's own accounting. The First Axiom of Management holds that the variety absorbed horizontally by the operational units must be balanced by the variety supplied vertically on the channels of cohesion. If AI raises the horizontal term — more capable units, cheaper coordination between them — then the same cohesion can be maintained with a *smaller* vertical term. The balance can be struck lower on the command axis. That is the structural reason AI can enable flatter, more networked organizations: not that coordination disappears, but that it no longer has to climb the hierarchy and descend again in order to happen.

The hinge of the whole movement is System 2. Recall that System 2 is the anti-oscillation function — the coordination that keeps autonomous System 1 units from colliding as they each optimize locally — and that Beer is insistent it does *not* sit on the command axis. Historically System 2 stayed thin, because rich, real-time, cross-unit coordination was too expensive to run, so organizations handled oscillation the cheap way: they pushed the conflict up into System 3 and let a manager command it away. AI is almost a native System 2 technology: in principle it can watch every unit's state and plans continuously, detect an emerging conflict, and propose or enact a harmonizing adjustment — coordination without a boss. (In principle, because deployed systems of this kind are still young; the claim is about the shape of the capability, not a report from the field.) Its deepest structural effect is therefore to let System 2 absorb load that used to be forced up onto System 3's command axis. That substitution — coordination for command — is precisely what lets an organization widen autonomy without dissolving into thrash. It also attacks Hoverstadt's Control Dilemma

at the root: managers micro-manage because coordination is failing and they compensate with command, so give them a System 2 that genuinely works and the compensatory hierarchy relaxes. A strong System 2, in short, is the enabling condition for distributed autonomy.

And here the argument turns back on itself, which is the point. The same capability that can strengthen System 2 into a coordination layer that *enlarges* autonomy can just as easily be pointed at System 3 and turned into a surveillance-and-command layer that *eliminates* it. The center that can now query raw state on demand can use that power to distribute decisions or to monitor them out of existence. The introduction called this "when the pen is a machine," and it is the modern, literal form of Beer's fault line: the identical cybernetic capability builds a liberty machine or an instrument of oppression, and the structure alone will not tell you which one you have built. Flatness is not freedom. A network can concentrate control as thoroughly as a pyramid, and do it behind a friendlier interface. We need a criterion that structure cannot supply.

The Criterion: The Agency Divide

The criterion is agency, and the reframing it demands is the one I have argued for elsewhere: the socially consequential question about AI is not whether the machine is autonomous, but how the deployment redistributes the human capacity to act. The prevailing debate asks how much agency humans keep versus how much the machine takes — as if it were a single quantity moving between two parties. That framing misses the mechanism. As I put it in the agency work, "AI does not redistribute agency by itself; institutions redistribute it *through* AI." The machine is the instrument; the design and the institution decide who gains and who loses. This is Beer's claim about the liberty machine, restated for a technology he never saw.

Agency, on this account, is not a single dial but a configuration of five dimensions, and naming them turns the vague word into something you can inspect. *Capacity* — can I act at all? — is knowledge, skill, tools, time. *Room* — may I, and can I, act *here*? — is rights, authority, access, the institutionally available space. *Will* — will I act? — is motivation, self-efficacy, ownership. *Direction* — toward what do I act? — is goals, values, sensemaking, the images of the future one is steering by. And *Identity* — who acts, and for what? — is the self- and role-conceptions from which responsibilities and purposes emerge. The dimensions are interdependent, not additive: Identity shapes which Directions feel fitting, Direction focuses Will, Will mobilizes Capacity, and Room decides whether mobilized capacity becomes effective action at all.

Set beside the Viable System Model, the mapping is close enough to be useful. *Room* is autonomy in Beer's exact sense — the residual variety a System 1 unit has to act within its resource bargain, the operational form of the "freedom to act" that *Designing Freedom* is about. *Identity* is System 5. *Direction* is the sensemaking work of Systems 4 and 5. And what the agency

model calls *meta-agency* — the institutional power to determine "who may revise rules, inspect data, select models, or challenge decisions" — is, in effect, *who holds the pen*, lifted from the single organization to the society around it. The set's central question turns out to be a special case of the agency divide.

The value of the five dimensions is that they sharpen every caveat the structural argument could only gesture at. A flat, AI-coordinated network sounds like more freedom, but the agency lens asks which dimensions actually moved. A recurring pattern is that AI raises *Capacity* while cutting *Room*: the worker can now do far more, and is watched far more closely doing it, with less ability to negotiate the target, the pace, or the rating. That is not the spread of agency; it is its concentration behind a more capable tool.

A composite scene — illustrative, like the cases of Part III — makes the pattern concrete. A parcel-delivery firm gives its dispatchers an AI routing assistant. Within a quarter each dispatcher handles half again as many vehicles, reads weather and traffic they never used to see, and reroutes in seconds: Capacity, up unmistakably. But the assistant's recommendations are logged; deviations must be justified in a comment field; the justifications feed the next performance review; and the target quietly ticks up to the new normal. No one on the floor may adjust the algorithm's priorities, and no one is quite sure who could. Room — the effective right to depart, the standing to negotiate the target, the access to the rule-making — has gone, without a single decision that said so. Every dispatcher is more capable; each is less free. Ask them who holds the pen and they will point, accurately, at a system nobody in the building can revise.

The agency work states the test bluntly: "an explainable system without an effective right to depart from its recommendation remains a system of control," and empowering people without giving them resources or authority merely turns "autonomy into responsibility without power." Both describe, in agency language, the pathology the VSM already knows — the System 2 that has quietly stopped harmonizing and started commanding, the coordination layer that has become covert System 3.

Read across the whole span, the same idea appears three times in three vocabularies, fifty years apart. Beer in 1973: freedom is designed, and the tools serve liberty or oppression depending on how you build them. The seven methods of this set: structure places authority over the model, and the placement is a choice. The agency work in 2026: the divide that matters is the distribution of human agency, and institutions do the distributing through AI. It is one argument. The chapter you are reading is the place the set says it in full.

Designing for Agency: From Diagnosis to a Check

A criterion you cannot apply is only a sermon, so the last move is from diagnosis to design. The agency work proposes an *Agency Impact Assessment* —

a companion to the risk and rights assessments that AI governance already runs — built to ask, before a deployment and repeatedly after it, the question structure cannot answer: which people and groups gain, and which lose, the durable ability to shape their own and shared futures? Its instruments are deliberately modest and cybernetic in spirit: a distributional analysis of who wins and loses across the five dimensions; an honest mapping of what is uncertain; reversible experiments rather than irreversible rollouts; and continuous learning from what the deployment actually does. It is, in effect, Beer's insistence that freedom be *designed* rather than assumed, turned into a procedure you can run.

The seven methods differ in how readily they carry such a check, and that difference is itself a useful guide. The participative methods serve agency distribution almost natively: Espinosa and Walker's self-transformation, where the members author their own model and own the indicators, and Espejo and Reyes's identity-first VIPLAN, where the naming is negotiated in the room, both build Room and Identity into the method rather than bolting them on. The expert and solo methods — Beer's own charts, worked alone; Pérez Ríos's analyst driving software — produce rigor and comparability that the participative methods cannot, but they concentrate authorship, and an agency check has to be added consciously or it will not happen. Hoverstadt's dynamics lens earns a special place here, because his whole method is trained to see where structure crushes the capacity to act — his Control Dilemma is an agency failure diagnosed in structural terms, a decade before the vocabulary existed.

The design stance that follows is simple to state and hard to hold. Build the System 2 layer to *harmonize, and not to command*: the moment an AI coordination system begins deciding rather than deconflicting, Room collapses and you have recentralized while calling it coordination. And keep the functions where identity and judgment live in human hands. The showcase on AI in this set makes the point in language worth keeping: the interpretive work of System 4 — what does this pattern *mean for us?* — and the trade-offs of System 5 — what would we never trade away? — are not tasks awaiting automation, because they are the functions that define the organization's viability and identity, and "delegating them is not automation, it is abdication." That is where the pen must stay in a human hand, not out of sentiment, but because handing it over is handing over the organization's capacity to remain itself.

Beer ended *Designing Freedom* with a warning and an invitation: that we were sleepwalking toward institutions no one would have chosen, and that the alternative was available to anyone willing to design for it. Half a century later the machine on the fault line is a better one, and the invitation is the same. The Viable System Model tells you how to build a system that survives; *Designing Freedom* and the agency divide tell you whom it should survive *for*. That second question is not one the model answers. It is a question the designer answers — and the honest first move, now as in 1973, is to admit that the question is being answered at all.

What to Borrow

Ask the freedom question first. Before adopting any structure or deploying any system, ask Beer's question directly: is this a liberty machine or an instrument of control — does it enlarge the room people have to act, or shrink it? The question sounds rhetorical and is not; most designs answer it by default, unexamined.

Use agency, not flatness, as the test. A flatter or more networked organization is not automatically a freer one. Measure the thing that matters — the distribution of Capacity, Room, Will, Direction and Identity across the people affected — and be especially alert to the common trap of raising Capacity while cutting Room.

Keep System 2 anti-oscillatory. AI makes System 2 powerful, and a powerful coordination layer is a liberating thing right up until it starts deciding. The moment coordination accretes decision rights it has become System 3 in disguise, and autonomy is gone. Watch that line; it is where the politics lives.

Never delegate System 4 interpretation or System 5 trade-offs. What a pattern means for us, and what we would never trade away, are the functions that constitute identity and viability. Automating the drafting beneath them is fine; handing over the judgments themselves is abdication, not efficiency.

Run an agency impact check before you deploy. Make the freedom question a procedure, not a sentiment: who gains and who loses agency, what is uncertain, can it be reversed, and what is it actually doing once live. It is *Designing Freedom* turned into something you can put on an agenda.

The Method Set

This chapter belongs to the set profiling seven VSM methods against a common template, divided by one question: *who holds the pen, and with what?* The annotated roster of the set is given once, in the opener to Part I; ch. 8 lays the seven side by side and turns the comparison into a decision guide.

References

- Beer, S. (1974). *Designing Freedom* (CBC Massey Lectures, 1973). Toronto: CBC Publications; Chichester: Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. ISBN 0-471-90675-1.

- Espejo, R. & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-19108-4.
- Espinosa, A. & Walker, J. (2017). *A Complexity Approach to Sustainability: Theory and Application* (2nd ed.). Series on Complexity Science, Vol. 5. Singapore: World Scientific. ISBN 978-1-78634-203-4.
- Hoverstadt, P. (2008). *The Fractal Organization: Creating Sustainable Organizations with the Viable System Model*. Chichester: John Wiley & Sons. ISBN 978-0-470-06056-8.
- Mathis, K. (2026). *Human Agency in the Age of Artificial Intelligence: From Controlling Technical Systems to Shaping Societal Capacity for Action*. Working paper, July 2026.
- Pérez Ríos, J. (2012). *Design and Diagnosis for Sustainable Organizations: The Viable System Method*. Berlin & Heidelberg: Springer-Verlag. ISBN 978-3-642-22317-4.

Part III

Part III — The Model at Work

Six Showcases: the VSM Mechanisms in Practice

The first two parts of this book are about *methods*. Part I profiles seven of them and holds each to the same question — who holds the pen? Part II turns that question into a decision. All of it operates one level up from the organization itself: it is about how you would go about using the Viable System Model, and who gets to decide. After the showcases, the coda that closes the book turns the same question into a thesis about freedom and agency.

Part III drops down a level. It is about the *mechanisms* — the parts of a viable system doing their work in an organization, on an ordinary Monday. The six showcases that follow each take a single piece of the model, put it under load in a concrete situation, and show what good design does to it. They are the machinery running, not the machinery described. One honesty note before the first of them: the organizations in these showcases are composites, drawn from consulting practice and the case literature — names are invented, and figures are illustrative, not reported. Where a real, published case is meant, it is cited as one.

They share a shape, because the mechanisms do. Each opens on a chronic tension that most organizations live with and few resolve — a today that crowds out tomorrow, a signal that dies on its way up, a bargain that curdles into either neglect or surveillance, teams that duplicate or collide across levels, a strategy no one can steer by, an argument about AI that has become a referendum. Each then does the move this whole book keeps returning to: it stops trying to declare a winner and designs the exchange instead. And each ends where the Tautai tradition ends — with a handful of Monday-morning actions and a navigator's parable, because a mechanism you cannot act on by Monday is a diagram, not a design.

The six:

- **Showcase 1 — Spanning the Adaptive Space**: the System 3–System 4 homeostat, refereed by System 5; balancing today's operations against tomorrow's adaptation without splitting the organization into two.
- **Showcase 2 — The Gatekeeper's Dilemma**: the System 1 → System 4 signal channel — a designed extension of canonical VSM — and the algedonic bypass; turning middle managers from signal filters into signal engineers.
- **Showcase 3 — The Resource Bargain**: the System 1–System 3 accountability loop, stabilized by System 2 and audited by System 3*; autonomy exchanged for accountability without collapsing into control.
- **Showcase 4 — The Republic of Teams**: recursion — every viable unit holding its own Systems 1–5 — and the designed dialogue between one level and the next.
- **Showcase 5 — Steering Without Shouting**: System 5 as variety attenuator; identity and narrative as the channel that aligns thousands of local decisions without issuing instructions.

- **Showcase 6 — Where the Machine Belongs**: placing AI across all five systems as a variety-allocation question, and drawing the judgment boundary that keeps interpretation and identity human — the practical companion to the coda.

You can read them straight through, after the methods, as a set of worked examples. Or you can ignore the order entirely and open the one whose tension you recognize from your own organization — which is, after all, how a navigator reads: not front to back, but starting from whatever is actually in the water in front of you.

The mechanisms in Part III are diagnosed and designed with the methods profiled in Part I and chosen among in Part II. Each showcase closes with its own sources and its companion chapters in The Tautai Principle.

Chapter 10

Spanning the Adaptive Space

Balancing Today's Operations Against Tomorrow's Adaptation — Without Splitting the Organization in Two

In Brief

Every organization has to do two incompatible things at once: run today's business as efficiently as it can, and change that business before the environment forces the issue. The two pull in opposite directions — one toward optimization and stability, the other toward disruption and renewal — and most organizations resolve the tension the way it feels natural to: they put the two in separate boxes. A core that executes, and, off to the side, an innovation lab, a transformation office, a strategy team that adapts. The "dual operating system," in John Kotter's phrase.

It disappoints more often than its popularity suggests, and the Viable System Model explains why. Running today and changing for tomorrow are not two functions to be housed apart. They are the two ends of a single balance that has to be struck continuously — in the same organization, over the same resources, again and again. In VSM terms this is the **System 3–System 4 homeostat**: System 3, the *inside-and-now*, optimizing what already exists; System 4, the *outside-and-then*, sensing what is coming; and System 5 holding the balance between them. The adaptive space is not a place, a unit, or a budget line. It is the managed tension between these two functions — and spanning it means *designing the exchange*, not separating the parties.

This showcase shows how. It names the mechanism, shows what happens when the balance is left to a wall between two departments, and gives the design that keeps the two in productive tension instead of mutual starvation.

The Innovation Lab Nobody Listened To

A mid-sized manufacturer felt the future arriving and did the responsible thing: it stood up a digital innovation unit. Bright people, a real runway, a corner of the building with better coffee and permission to think three years out. Two years later the unit had produced a wall of elegant decks, a handful of pilots, and almost no change to the business that pays the bills.

Nobody involved did anything wrong. The core was measured, quarter by quarter, on cost and throughput; absorbing an unproven idea from the lab was, for any plant manager, a threat to the numbers they answered for. The lab, for its part, was cut off from the operational reality that would have made its ideas land — it was dreaming without friction. So the two sides waved at each other across a wall and handed off proposals that died on contact. The organization had built two rooms and told them to cooperate by memo.

The opposite failure is quieter and more common: the organization with no meaningful System 4 at all. Superb at execution, relentlessly efficient — and repeatedly surprised by changes everyone outside could see coming. It optimizes beautifully, right up to the moment it discovers it has been getting better and better at something the market no longer wants. It marches, in good order, into irrelevance.

Both failures are the same failure seen from two sides: an organization that has not designed the place where today and tomorrow are supposed to negotiate.

Why It Is a Homeostat, Not a Department

Here is the reframe. System 3 asks one question — *how do we optimize what we are already doing?* System 4 asks the opposite — *what should we be doing differently?* The tension between present efficiency and future adaptation is not a problem to be solved once and filed away. It is structural and permanent. It has to be *held*, continuously, the way a body holds its temperature — which is why Beer called the arrangement a homeostat: a self-balancing loop between two forces that must never let either win outright.

Let either win and the organization dies in a recognizable way. When System 3 dominates, adaptation gets starved to feed optimization, and the business perfects itself into obsolescence. When System 4 floats free, it builds what practitioners call castles in the air — strategies and innovations with no operational grip, the wall of decks the manufacturer's lab produced. Neither is a villain. Each is one half of a balance that has lost its other half.

This is why System 5 exists. Its first responsibility, before identity or policy or anything else, is the one the textbooks list first: *balancing Systems 3 and 4*. System 5 sets where the balance point sits between running today and building tomorrow, and owns the trade-off when the two cannot settle it themselves — a trade-off it should make deliberately and sparingly, intervening only when the homeostat below it genuinely cannot resolve the matter.

And now the diagnosis of the dual operating system writes itself. Splitting System 3 and System 4 into separate organizational units does not remove the homeostat — nothing can, because the tension is intrinsic to being a viable system. All the split does is sever the channel along which the balance is supposed to be struck, and hand the balancing act to a turf

war between two departments that neither can win. The innovation lab was not a bad idea badly executed. It was a homeostat with its two halves put in different rooms.

One more thing the model insists on: the balance is struck at *every* level, not only at the top. Every viable unit — each division, each team — runs its own small System 3 and System 4, and therefore has its own adaptive space to span. Spanning it is not the chief executive's private task. It is a discipline that has to live at every recursion, or it does not really live at all.

The Design: A Balanced, Refereed Exchange

If the adaptive space is a homeostat, the design is not a lab. It is a rhythm — a standing, resourced conversation between the people who run today and the people who read tomorrow, refereed by the people who own the balance. Four moves make it real.

Make the balance a visible number. The split between running-today and changing-tomorrow — of budget, of attention, of the best people's time — should be an explicit allocation decision, made on purpose at System 5 and revisited on a rhythm. The failure mode it prevents is the default one, in which "innovation gets whatever is left after operations has taken what it needs," which is to say nothing, reliably.

Convene the exchange, and let each side attack the other. System 4 brings the signals and scenarios; System 3 brings operational reality; and the point of the meeting is friction, not harmony. System 3's job is to say *that will not survive contact with the plant floor*; System 4's job is to say *that efficiency gain quietly forecloses the pivot we are going to need*. Out of that grinding comes what neither produces alone: strategy grounded in what the organization can actually do, and operations that understand why they are being asked to change.

Let System 5 hold the balance point, and only that. The executive's role is not to attend every exchange or adjudicate every idea. It is to set where the balance sits, to guard the boundaries, and to own the genuine trade-off when the homeostat deadlocks — and then to step back. That trade-off, incidentally, is exactly the constitutive judgment the AI showcase (Showcase 6) argues you must never automate: deciding how much of today to spend on tomorrow *is* the organization choosing what it will become.

Keep it inside the line, at every level. Because the balance is struck at each recursion, the design is not one grand session at the top but a small, repeating S3–S4 conversation owned inside each unit — so that adaptation is something the operating business does, not something done to it by a separate office.

Seen in Practice: The Firm That Dissolved Its Transformation Office

The manufacturer's wall is not peculiar to manufacturing. A regional logistics company ran the same arc with a central transformation office instead

of a lab: two years of roadmaps the depots ignored, because every pilot hour was an hour lost against the on-time metric. The redesign dissolved the office outright. In its place, each region got a standing monthly session in which the operations lead and a market-facing scout jointly owned a small, ring-fenced adaptation budget, its size set quarterly by the executive team. Within two quarters, pilots that had died on contact were reaching live operations — not because the ideas got better, but because the depot that would have to run each change now sat in the room while it was shaped.

Failure Modes to Design Against

The headless chicken (no System 4). Efficient, blind, and repeatedly surprised. If nobody is resourced to read the horizon, the organization will optimize its way into a future it never saw coming. The fix is not a slogan about innovation; it is a funded, listened-to sensing function.

The dominant System 3. Optimization crowds adaptation out of the calendar and the budget until "we'll get to that after this quarter" becomes the permanent condition. Watch for an adaptation budget that is always the first thing cut.

Castles in the air (detached System 4). Strategy and innovation with no operational grip — the wall of decks. If the sensing function has no channel into operations and no operational co-owner, its output is entertainment.

The dual-operating-system split. Housing System 3 and System 4 in separate units and hoping a hand-off substitutes for a homeostat. The tension does not go away when you separate the parties; it just stops being productive.

System 5 abdicates the balance. When leadership declines to set the balance point, the split defaults to whoever shouts loudest or to the residue of the budget cycle — and it drifts, every time, toward the short term. Holding the balance is not a distraction from the executive job. It is the executive job.

Which method finds this. Beer's diagnosis catches this wall directly: Chart Four of his Three–Four homeostat asks who actually meets whom, and with what shared model — and here the answer is nobody, by memo (ch. 1). Hoverstadt's pathology bestiary names both ends of the imbalance: Bean Counters when System 3 dominates, Castles in the Air when System 4 floats free (ch. 7). Schwaninger's three logical levels explain the earliest symptom — an organization winning on the metrics and losing somehow anyway, because profit is being managed while viability is not (ch. 6). The pen here belongs to System 5: whoever sets the balance point holds it, and the redesign only sticks if they hold it on purpose.

Monday Morning Actions

- **Trace one strategic insight to an operational change.** Take something your sensing function saw in the last year and follow it. Did it

change how the business actually runs? If not, the homeostat has a broken channel — find where it broke. Owner: the executive sponsor.
- **Make the balance a number.** State, explicitly, the share of resource and attention going to run-today versus change-tomorrow, and decide it on purpose rather than by default. Owner: CFO with the executive team (System 5).
- **Convene one real S3–S4 exchange.** A standing session where an operations lead and a strategy/market lead jointly own one adaptation decision and a protected budget for it. Owner: COO with Head of Strategy.
- **If you have a separate innovation unit, give every project a core co-owner.** No pilot proceeds without someone from the operating business who will have to run it, in the room while it is shaped. Owner: the innovation unit's sponsor.

The apprentice asked: "Master, the paddlers want to drive straight and fast, and you sit reading clouds that are not yet weather. Whom should the captain believe?" The master navigator smiled. "Both, always, and neither alone. The paddlers keep us moving now; without them we drift and starve. But their rhythm is loud, and it would gladly drown the news of a storm they cannot yet see, or an island not yet above the water. My task is to hear what is not yet here. The captain's task is the hardest of the three: to let the paddlers drive without letting their drumbeat deafen the canoe to tomorrow — and to let me dream without letting the dreaming stop the paddles. A vaka that only paddles arrives swiftly at the wrong place. A vaka that only reads the horizon never leaves the lagoon."

Sources & Related Chapters

The System 3–System 4 homeostat and System 5's balancing role are diagnosed and designed with the methods profiled in Parts I and II:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — the Three-Four homeostat, System 5 as the balancer, and the "inside-and-now" / "outside-and-then" distinction, in the originator's own worksheets.
- Pfiffner, *The Neurology of Business* (ch. 2) — the S3–S4 balance as a design step, and System 5's minimal-intervention role.
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — the S3–S4 balance check and adaptation as a whole-system capacity.
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — the "headless chicken" (absent System 4) and the S4–S3 dissociation pathologies, catalogued.

- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management and viability *beyond* survival), and **Hoverstadt** (ch. 7; the pathology bestiary — "Bean Counters" and "Castles in the Air" as the two ends of this imbalance).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- For the set's thesis on AI, System 2 and human agency, see **Designing Freedom, Again** (ch. 9) — the coda.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1981). *Brain of the Firm* (2nd ed.). Chichester: Wiley.

Companion Tautai chapters: Chapter 7 (Strategic Adaptivity); Chapter 15 (Spanning the Adaptive Space).

Chapter 11

The Gatekeeper's Dilemma

Turn Your Middle Managers from Signal Filters into Signal Engineers

In Brief

Every organization asks its middle managers to do two contradictory things at once: protect the organization from information overload by filtering what travels upward, and ensure that no critical market signal ever gets lost on the way. Most organizations resolve this contradiction silently — in favor of filtering. The result is a transmission gap: world-class sensing at the edge, strategic blindness at the center.

This showcase shows how to resolve the contradiction structurally rather than heroically. The Viable System Model provides the design logic: not one contested reporting channel, but two designed channels with different rules — a managed routine channel where filtering is legitimate and valuable, and a weak-signal express lane from operations (VSM System 1) to strategic intelligence (VSM System 4), supplemented by an emergency bypass for pain signals. The conflict between managers who filter and strategists who need raw signals becomes a productive exchange about signal quality — and the middle manager's role converts from gatekeeper to signal engineer.

The Manager Who Did Everything Right

Consider a mid-sized logistics firm (readers of The Tautai Principle will recognize it from Chapter 5 there). A frontline customer success team sensed a clear market opportunity — a simplified, all-in-one customs clearance product their SME clients kept asking for. They wrote it up. They sent it to their regional manager. He archived it. Eighteen months later a competitor launched exactly that product and took a substantial share of the firm's SME market.

It is tempting to make the manager the villain of this story. Resist the temptation, because here is the uncomfortable truth: **he was doing his job.**

His job, as the organization had defined it, was to keep his region focused on quarterly efficiency targets. His job was to shield his teams from distraction and his superiors from noise. Every incentive, every metric, every unspoken norm told him that a speculative product proposal from

a customer success team was exactly the kind of thing a good manager absorbs so that the people above him can concentrate on what matters.

The organization did not suffer from a bad manager. It suffered from a bad design — a design that forced one human being to be simultaneously a filter and a conduit, with no rules for deciding which signal deserves which treatment.

This is the Gatekeeper's Dilemma. You cannot solve it by exhorting managers to "let more through." You solve it by giving different kinds of signals different pathways.

Why Filtering Is a Feature, Not a Betrayal

Before we redesign anything, we must rehabilitate the filter.

A cybernetic truth, going back to Ross Ashby's Law of Requisite Variety, sits underneath every organization chart: the environment always contains vastly more variety — more events, more signals, more detail — than any management team can absorb. If every customer complaint, every competitor rumor, and every process hiccup traveled to the executive floor unfiltered, the strategic brain would drown in noise within a week. Attenuation — the deliberate reduction of variety on the way up — is not a pathology. It is a survival requirement.

The middle manager who summarizes, prioritizes, and quietly deletes is performing variety engineering. The organization needs this function desperately.

The pathology begins when attenuation is *undesigned* — when one channel must carry everything, and one person's judgment, shaped by operational incentives, silently decides what the strategy function will never see. Then two legitimate perspectives collide:

- **Operational management says:** "Strategy drowns us in requests and drowns themselves in noise. My job is to protect focus."
- **Strategic intelligence says:** "Middle management buries our most valuable insights. By the time information reaches us, it has been sanitized into uselessness."

Both are right. Both are describing the same broken design from opposite ends. And as long as there is only one channel, the fight over what travels through it is a zero-sum political battle, in which a manager's power derives from being the node through which communication must pass — the politics of connection, as The Tautai Principle (ch. 5) calls it.

The adaptive move is the same one the previous showcase made with the System 3–4 conflict: stop trying to declare a winner, and design the exchange instead.

The VSM Design: Three Channels, Three Rule Sets

The Viable System Model does not run a single pipe from bottom to top. It works with distinct channels with distinct purposes. Applied to the Gatekeeper's Dilemma, three channels matter — two of them canonical Beer, one an extension this book proposes in the same spirit. The section on Channel 2 says exactly where the extension departs from the original model.

Channel 1: The Managed Reporting Line (System 1 → System 3)

This is the classic vertical channel: performance data, resource needs, operational exceptions flowing from value-creating teams to operational management. Here, filtering is not just permitted — it is the point. System 3 needs comparable, aggregated, prioritized information to run the current business. The manager who condenses twelve team reports into one page of decisions-needed is creating value.

Rule set: attenuate aggressively. Summarize, standardize, prioritize.

Channel 2: The Weak-Signal Express Lane (System 1 → System 4)

This is the channel most organizations never build. Frontline teams are your organization's sensory surface — they touch customers, suppliers, and competitors daily. The patterns they notice are the raw material of strategy. But weak signals, by definition, do not look important yet. They fail every prioritization test that Channel 1 rightly applies. A signal that survives aggressive attenuation is no longer weak — which means Channel 1 structurally cannot carry the early signals System 4 needs most.

So the express lane runs *around* the operational hierarchy, not through it. Any frontline team can file a market observation directly into the strategic intelligence function — lightweight, structured, unfiltered. No manager approval required. No manager blame attached.

Rule set: attenuate lightly, and never on grounds of operational convenience. Structure the input; do not judge its strategic worth at the point of origin — that judgment belongs to System 4, which sees the pattern across many signals.

A word to readers who know their Beer: this channel is not in the original diagrams. Beer's vertical channels all run between the operational units and System 3; System 4 connects to the environment and to System 3, and the only canonical bypass from operations upward is the algedonic channel to System 5. The express lane is an extension — but one the model itself argues for. Each System 1's local environment is embedded in the total environment System 4 must scan, which makes frontline teams System 4's cheapest and best-placed sensors; the express lane merely gives that embedding a designed pathway.

The canonical element the express lane will remind such readers of is System 3* — the sporadic audit channel, which also bypasses the System 1 management filter to obtain unfiltered ground truth (Showcase 3 treats it in depth). The resemblance ends at the bypass. System 3* is initiated top-down by System 3, works sporadically, examines internal operational reality — the inside-and-now — and serves control: it is what makes the

resource bargain trustworthy. The express lane is initiated bottom-up by the teams themselves, runs continuously, carries external environmental observations — the outside-and-then — and serves intelligence. An audit of operations, however thorough, will never surface a weak market signal. That is why 3* cannot substitute for this channel, and why the extension is needed at all.

Channel 3: The Algedonic Bypass (Anywhere → System 5)

Stafford Beer reserved a special channel for signals of acute pain or opportunity — *algedonic* signals, from the Greek for pain and pleasure. These are the signals that must not wait for weekly syncs or monthly reviews: the safety risk, the reputational time bomb, the customer defection wave that has already started. The algedonic channel bypasses every filter and alerts the top of the organization directly, the organizational equivalent of a hand jerking away from a hot stove before the conscious brain knows why.

Rule set: no attenuation at all — but strict admission criteria, so the channel stays credible. An algedonic channel that carries routine complaints is dead within a quarter.

Notice what this three-channel design does to the political conflict. The manager is no longer the single point through which everything must pass — so the fight about what he "should have" forwarded dissolves. Filtering on Channel 1 stops being a betrayal of strategy. Direct signals on Channel 2 stop being an act of insubordination. The question shifts from *who controls the channel* to *is each channel healthy* — and that is a question both sides can work on together.

The Productive Exchange: From Turf War to Signal Quality Dialogue

With the channels separated, the conflict between operational management and strategic intelligence can finally become what the adaptive space (Showcase 1) requires: a structured, recurring, productive exchange. It unfolds in a rhythm much like the System 3–4 collaboration process:

Step 1: Signal Review. System 4 regularly shares back what arrived through the express lane — which signals formed patterns, which triggered exploration, which came to nothing. Frontline teams and their managers see, for the first time, what happens to what they send.

Step 2: Calibration. Operational managers challenge the signal flow from their side: "These three signal types are pure noise — here is why." System 4 challenges back: "This category you consider noise predicted the last two market shifts." Together they tune the structured input format — not by adding approval steps, but by improving the questions the express lane asks.

Step 3: Sensing Priorities. System 4 tells the front line what it is currently trying to see: "We need early indicators on customer reaction to the new regulation." Sensing becomes directed without becoming censored. This is the scanning system of The Tautai Principle (ch. 4) acquiring a steering wheel.

Step 4: Feedback on Consequence. Nothing kills a sensing culture faster than silence. Every signal that enters the express lane receives a disposition — pattern-matched, parked, escalated, retired — visible to its originator. The logistics team that proposed the customs product would have known, within weeks, that their insight was alive.

Step 5: Role Recognition. Managers are measured and recognized for the signal performance of their area: how fast signals travel, how many of their teams' observations contribute to strategic patterns, how well they coach teams to sense. The Signal Velocity metric — time from frontline detection to strategic review — becomes a management KPI, not just a system health indicator.

The expected outcomes mirror the System 3–4 exchange of Showcase 1: grounded strategy (built on live frontline evidence rather than sanitized reports), operational teams that understand *why* their observations matter, shared intelligence across the operational–strategic boundary, and a middle management layer whose power now grows with the quality of flow rather than the control of flow.

The Redefined Role: Signal Engineer

The manager's role must shift from gatekeeper of information to capability builder and enabler of flow — the argument The Tautai Principle (ch. 5) makes at length. The three-channel design makes that shift concrete. The signal engineer:

- **Runs Channel 1 with pride.** Aggregation and prioritization for operational control remain a core craft. Nothing in this design diminishes it.
- **Protects Channel 2 without touching it.** The manager ensures teams have the time, tools, and psychological safety to file signals — and conspicuously refrains from pre-approving them. The measure of success is flow, not control.
- **Coaches sensing as a skill.** Which observations are worth capturing? How do you describe a pattern instead of an anecdote? The manager teaches the difference — this is the capability-builder role in practice.
- **Guards the algedonic threshold.** The manager helps teams calibrate what genuinely belongs on the emergency channel, keeping it rare and therefore credible.

This is not a diminished role. It is a harder and more valuable one. Any organization can find people to summarize reports. Few have managers who can raise the sensing quality of an entire region.

Failure Modes to Design Against

The express lane becomes a complaint box. If Channel 2 fills with grievances about internal processes, it will be ignored. Route internal

improvement ideas to your improvement process; reserve the express lane for signals about markets, customers, competitors, and environment.

System 4 goes silent. If signals disappear into the strategy function without feedback, the front line stops sending within two quarters. The feedback loop of Step 4 is not a courtesy — it is the load-bearing wall.

Managers are punished by surprise. If an executive reads a frontline signal and calls the manager to ask why he didn't know about it, you have just re-installed the gatekeeper — every manager will demand pre-review. The protocol must be explicit: signals on Channel 2 carry no manager accountability. Surprise is the design working.

The algedonic channel inflates. Every escalation that turns out to be routine erodes the channel. System 5 must visibly triage: genuine algedonic signals get immediate response; misfiled ones get redirected with a short explanation of the criteria — without punishing the sender, or the channel chills.

Measuring the Health of the Channels

Metric	Channel	What it tells you
Signal Velocity: median time from frontline observation to System 4 review	Express lane	Whether intelligence travels at market speed
Feedback Latency: median time from signal filed to disposition visible	Express lane	Whether the sensing culture will survive
Signal Yield: share of strategic initiatives traceable to frontline signals	Express lane	Whether the channel matters, not just exists
Attenuation Quality: System 3 satisfaction with routine reporting	Reporting line	Whether filtering still serves operations
Algedonic Precision: share of bypass signals confirmed as critical	Bypass	Whether the emergency channel stays credible

Seen in Practice: The Bank That Filed Its Own Signal

The logistics firm's arc repeats in banking, beat for beat: branch managers at a regional retail bank noticed SME clients moving their foreign-exchange business to a fintech app, area managers — measured on deposit growth and cost-income ratio — filtered the observation as noise, and a neobank

took the segment. The repair was the express lane in miniature: a one-field *market signal* form flowing from any branch straight to the payments-strategy team, no sign-off, no blame, dispositions posted back monthly. Within a quarter the FX pattern arrived again, from several regions at once, and a same-day FX product entered the roadmap. What had changed was not managerial diligence but signal velocity — the first metric in the table above, moved on purpose.

Which method finds this. Beer's diagnosis starts by walking the six vertical channels: trace the dead signal along them and you find a channel that was never built, and no algedonic route around it (ch. 1). Hoverstadt's pathology bestiary files the same finding under the Missing Link, with the manager caught in his Control Dilemma — accountable for everything that passes upward, and so incented to pass less (ch. 7). Espinosa and Walker add the test the redesign must meet: participation only counts when it carries decision rights, and a suggestion box managers may ignore is consult-and-ignore, not sensing (ch. 3). In this scenario the pen sits with whoever defines the channels — until the express lane exists, the regional manager holds it by default, one archived memo at a time.

Monday Morning Actions

- **Trace one dead signal.** Find one market insight from the last year that originated at the front line and died en route. Map where it stalled and which channel it was forced to travel. Owner: Head of Strategy.
- **Open a minimal express lane.** One structured form — observation, evidence, pattern guess — flowing directly to the strategy function. Announce the no-blame, no-approval rule explicitly. Owner: Head of Strategy, endorsed by the COO.
- **Close the loop publicly.** Commit to a monthly signal review where dispositions are shared back to originators. Put the first one on the calendar before announcing the channel. Owner: Head of Strategy.
- **Rewrite one job expectation.** Add signal performance to the role profile of one management layer, and remove any implied accountability for pre-clearing team signals. Owner: COO with HR.

The apprentice asked: "Master, the lookout on the outrigger shouts many things — birds, driftwood, cloud shapes. Should the crew chief not decide which shouts reach you?" The master navigator shook her head. "The crew chief steers the paddlers, and for that he must quiet the noise. But the shape of a cloud that means land — that must reach me as it was seen, not as he judged it. So we have two calls on this vaka: the working call, which the chief may soften, and the far-seeing call, which no one may touch. A crew that confuses them paddles hard toward nothing."

Sources & Related Chapters

The channels, algedonic bypass and audit function used here are diagnosed and designed with the methods profiled in the VSM method-chapter set:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — the algedonic signal, the six vertical channels, and System 3* audit, in the originator's own worksheets.
- Pfiffner, *The Neurology of Business* (ch. 2) — Step VI on communication channels (capacity, intelligibility, timeliness) and the algedonic check.
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — real-time, exception-based algedonic monitoring in a self-governance setting.
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — communication-channel pathologies and the missing-algedonic-channel diagnosis (Ch. 3, §3.3).
- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management), and **Hoverstadt** (ch. 7; the pathology bestiary).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- For the set's thesis on AI, System 2 and human agency, see **Designing Freedom, Again** (ch. 9) — the coda.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1981). *Brain of the Firm* (2nd ed.). Chichester: Wiley.

Companion Tautai chapters: Chapter 5 (The Organizational Nervous System); Chapter 15 (Spanning the Adaptive Space).

Chapter 12

The Resource Bargain

How to Grant Real Autonomy Without Losing the Organization

In Brief

Dynamic Teams only create competitive advantage when they have genuine autonomy — and genuine autonomy is precisely what operational management cannot responsibly hand out for free. Most organizations oscillate between two failure modes: they re-centralize after the first incident (killing speed), or they let empowered teams drift apart (killing coherence). Both responses treat autonomy as a quantity to be increased or decreased.

The Viable System Model offers a third path: autonomy as the outcome of a *negotiation*. Stafford Beer called it the resource bargain — teams receive resources and freedom to act in exchange for explicitly agreed accountability, verified not by continuous supervision but by sporadic, announced audits. The chronic tug-of-war between empowered teams and operational management becomes a recurring, structured exchange with a known format. This showcase shows how to run that negotiation, how System 2 keeps the resulting autonomy from producing chaos, and why the audit function — properly designed — is the strongest protector of team freedom your organization has.

The Pendulum That Never Stops

A European industrial services company — call it Nordwerk — did everything the Dynamic Teams playbook of *The Tautai Principle* (ch. 3) recommends. It dissolved parts of its rigid hierarchy into Dynamic Teams: small, cross-functional units forming around emerging customer needs. For six months, the results were exactly as promised. Response times dropped. Frontline innovation flourished. One team invented a maintenance-contract variant that became the fastest-growing offer in the portfolio.

Then a team in the Rotterdam region made a pricing decision that lost the company a six-figure sum.

You know what happened next, because it happens everywhere. The CFO demanded controls. A new approval matrix appeared. Every discount beyond a token percentage now required sign-off from two levels up. Within a quarter, the Dynamic Teams were dynamic in name only — and the

maintenance-contract innovators quietly noted that their next idea would not be worth the paperwork.

Eighteen months later, with growth stalling, a new COO swung the pendulum back: "We trust our teams." Controls were loosened. Until the next incident.

This pendulum — empowerment, incident, clampdown, stagnation, re-empowerment — is one of the most reliable patterns in organizational life. It persists because both sides of the argument are correct. Teams genuinely need freedom to act on local information at local speed; that is the entire competitive logic of the Structured Network Organization. And operational management genuinely carries accountability for the whole; a leader who cannot answer for the aggregate results of thirty autonomous teams is not empowering — she is abdicating.

The pendulum is what you get when an organization treats autonomy as a *setting* — a dial someone turns up or down. The Viable System Model treats it as something else entirely: a *contract*.

The Resource Bargain: Autonomy as a Negotiated Contract

In the VSM, the relationship between operational management (System 3) and the value-creating teams (System 1) is not supervision. It is an exchange with obligations on both sides, which Beer called the resource bargain:

System 3 grants: resources — budget, people, tools — and a defined space of discretion: the decisions the team may take without asking anyone.

System 1 grants in return: explicit accountability — agreed outcomes, agreed boundaries, and the obligation to keep its own house in order within them.

Three properties make this a bargain rather than a delegation memo.

It is negotiated, not imposed. The team argues for the discretion it needs, based on what it senses in its part of the market; management argues for the accountability the whole requires. Both sides can demand things of each other. A team that wants pricing freedom must articulate how it will bound its own risk. A management that wants tighter reporting must show what decision the reporting serves. The negotiation itself — conducted well — transfers exactly the operational and strategic context each side lacks. It is the productive conflict between operations and strategy that *The Tautai Principle* stages between Systems 3 and 4 (ch. 15), one recursion level down: the same collision of local insight with systemic concern, resolved through exchange instead of verdict.

It is explicit. The pendulum thrives on ambiguity. When nobody knows precisely where team discretion ends, every incident becomes an argument about what "should have been" escalated — and every argument ends with a new blanket control. A written bargain converts that argument into a boundary review: *was this inside the agreed space or not?* If inside, the outcome is a learning input, not a governance failure. If outside, the issue is the specific boundary, not autonomy as a concept. The Rotterdam pricing

incident, under a bargain, triggers a renegotiation of the pricing boundary for one team — not an approval matrix for thirty.

It is renegotiated on a rhythm. Markets shift; teams mature; boundaries age. The bargain is revisited on a fixed cadence — quarterly works for most organizations — and, importantly, on trigger events: a boundary violation, a market shift detected by the Scanning System, a team's demonstrated track record earning wider discretion. Autonomy becomes something a team *grows into* by evidence, not something granted in a moment of cultural enthusiasm and revoked in a moment of fear.

System 3*: Why Sporadic Audit Protects Autonomy

The most counterintuitive element of the VSM design is the audit channel, which Beer labeled System 3* (three-star). It is the mechanism that makes the bargain trustworthy without making it suffocating.

Here is the problem it solves. If System 3 can only verify team accountability through continuous reporting, then reporting inevitably swells — dashboards multiply, review meetings lengthen, and continuous monitoring quietly reinstalls the supervision that autonomy was supposed to replace. Trust that must be verified continuously is not trust; it is surveillance with extra steps.

System 3* takes a different route: occasional, direct, announced inspection of the actual operation — not the team's reports about itself. A few times a year, someone from the operational management sphere looks directly at how a team really works: its books, its processes, its safety practices, its customer commitments.

Done in the audit spirit the VSM intends, this inverts the usual emotional loading of the word "audit":

- **Because verification is sporadic and direct, routine reporting can shrink.** The team spends its days working, not reporting. The bargain specifies the few numbers System 3 actually needs; System 3* covers the rest.
- **Because the audit examines reality rather than reports, gaming loses most of its value.** An announced audit can still be prepared for; what sporadic, in-depth inspection does is raise the cost of sustained pretense well above the cost of simply keeping the house in order.
- **Because the audit is announced and procedural, it is impersonal.** It is not triggered by suspicion; it lands on everyone eventually. Being audited carries no stigma — which is precisely what makes findings discussable.

Tell your teams the truth about this design: *the audit is the price of your freedom, and it is a low price.* The alternative to three audit days a year is 250 days of continuous reporting. Organizations that understand this defend their audit function as fiercely as their autonomy — because structurally, they are the same thing.

System 2: The Quiet Stabilizer

A bargain between System 3 and each team still leaves a gap: teams can be individually accountable and collectively chaotic. Thirty teams, each within its agreed boundaries, can still collide — two teams courting the same customer with different prices, incompatible tooling decisions, oscillating demand on shared resources.

The reflex response is to route these collisions upward: let management coordinate. The VSM's answer is System 2 — coordination *between* teams that does not pass through the command channel at all. Shared standards, common protocols, synchronized rhythms, mutual visibility: the anti-oscillation machinery Beer built into the model (ch. 1).

In the context of the resource bargain, System 2 has a specific job: **it absorbs the friction that would otherwise become an argument for reducing autonomy.** Every cross-team collision that gets resolved by shared protocol is a collision that never lands on a manager's desk as evidence that "the teams can't handle freedom." A weak System 2 therefore doesn't just cost efficiency — it actively feeds the clampdown side of the pendulum.

Practical System 2 mechanisms for a Dynamic Team environment include a shared customer registry with engagement rules (who is talking to whom), common definitions of done and quality floors, a resource marketplace with visible queues instead of escalated resource fights, and cross-team sync rhythms owned by the teams themselves. None of these are commands. All of them are variety absorbers.

Running the Negotiation: A Five-Step Format

Like the System 3–4 collaboration in *The Tautai Principle* (ch. 15), the resource bargain works best as a structured, repeating process rather than a heroic one-off conversation.

Step 1: The Team's Bid. The team states what it intends to achieve in the coming period, what resources it needs, and — crucially — what discretion it needs: the specific decision rights that would let it move at market speed. Each requested right comes with the team's own proposal for bounding its risk.

Step 2: The System View. Operational management responds with the whole-system picture: interdependencies the team cannot see, aggregate risk positions, resource constraints across all teams, non-negotiable compliance boundaries. This is not a rejection round; it is context transfer.

Step 3: Joint Boundary Drafting. Both sides draft the discretion space together: decisions the team takes alone, decisions requiring consultation, decisions requiring consent, and hard limits. The most productive question in this step is management's: *"What would you need to demonstrate for this consultation right to become an autonomous right?"* — because it converts the boundary from a wall into a development path.

Step 4: Accountability Specification. Agree the few outcomes the team answers for, the handful of indicators System 3 genuinely needs continuously, and what System 3* will look at when it visits. Explicitly delete reporting that serves no named decision.

Step 5: Commit and Publish. The bargain is written, short, and visible to neighboring teams — because published bargains are also System 2 infrastructure: they let teams anticipate each other's decision spaces without asking.

Between negotiations, the bargain runs. Incidents inside the boundaries are processed as learning by the team's own improvement loops (*The Tautai Principle*, ch. 6). Boundary violations trigger a review of that boundary with that team. And demonstrated performance triggers Step 3's development path: discretion widens on evidence.

What Changes for Leaders

The resource bargain redefines the operational leadership craft in the same way *The Tautai Principle*'s ch. 15 redefines the strategy craft. The leader of Dynamic Teams stops being an approver of decisions and becomes a *designer of decision spaces*. Her skill shows not in the quality of her interventions but in the rarity of their necessity: teams whose boundaries are well-negotiated, whose collisions are absorbed by System 2, and whose accountability is verified by a credible System 3* simply do not generate many moments that need her.

This is measurable. Track the share of team decisions taken without escalation that remain aligned with organizational intent — the Strategic Alignment indicator from *The Tautai Principle* (ch. 5). Track boundary stability: how often bargains are violated versus renegotiated on rhythm. Track the pendulum itself: count the blanket controls introduced after incidents. In a healthy bargain system, that last number trends toward zero — incidents produce boundary adjustments for one team, not control regimes for all.

Seen in Practice: The Hospital That Stopped Swinging

Nordwerk's pendulum swings in clinics as reliably as in factories. A hospital group gave its clinical departments real budget autonomy, clamped down to central sign-off after one department's procurement breach, and then — instead of swinging back on a change of mood — ran the five-step bargain with each unit: discretion spaces drafted jointly, a quarterly clinical audit (System 3*) replacing blanket pre-approval, and a second overspend triggering a boundary review for that one unit rather than a control regime for all. What the clinic adds to the factory story is the boundary type: some limits — patient safety, procurement law — are not negotiable, and naming them as hard boundaries is exactly what frees everything else for the bargain.

Which method finds this. Beer's diagnosis goes straight to the accountability channel: does a resource bargain exist, and is it backed by a System 3* audit channel rather than by swelling reports (ch. 1)? Hoverstadt reads the same pendulum as a broken resource-performance bargain and prescribes trust built by monitoring — discretion widening as the record justifies it (ch. 7). Espinosa and Walker add the decisive nuance from the Suma co-op, where a proposed monitoring system was rejected as "Big Brother" until each work group chose its own indicators (pp. 201–202): who chooses the measures decides whether the same instrument is surveillance or self-regulation (ch. 3). In this scenario the pen sits on the negotiating table — the team drafts its own accountabilities, and management countersigns the boundary.

Monday Morning Actions

- **Audit your last clampdown.** Take the most recent control introduced after an incident. Was the triggering decision inside or outside any explicitly agreed team boundary? If no boundary existed, the control is treating a design gap as a trust problem. Owner: COO.
- **Write one bargain.** Pick one Dynamic Team and run the five-step negotiation to a one-page written bargain: outcomes, discretion space, minimal reporting, audit scope. Treat it as a pilot for the format. Owner: the team's operational leader.
- **Delete orphan reporting.** For one team, list every recurring report it produces and name the decision each one serves. Delete the orphans, and announce that System 3* visits replace them. Owner: COO.
- **Charter one System 2 mechanism.** Identify the most frequent cross-team collision of the last quarter and have the affected teams design a shared protocol for it — without management in the room for the design. Owner: the senior leader both teams report to.

The apprentice asked: "Master, you let each paddler set his own stroke — yet the vaka does not scatter. How?" The master navigator held up the rigging. "Before we left shore, each crew member and I agreed what is his alone to decide and what belongs to the vaka. The drummer keeps us from fouling each other's strokes — that is not command, that is rhythm. And now and then I walk the hull and look at every lashing with my own eyes. Not because I distrust the crew — but because I check the lashings, I never need to watch their hands."

Sources & Related Chapters

The resource bargain, System 2 stabilization and System 3* audit used here are diagnosed and designed with the methods profiled in the VSM method-chapter set:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — the resource bargain and accountability channel, and System 3* as the sporadic audit that grants autonomy its cover.
- Pfiffner, *The Neurology of Business* (ch. 2) — Step IV (central vs. decentral) and the "Powers Reserved" principle: what is not forbidden is permitted.
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — devolved control and indicator ownership (the Suma "Big Brother" case, pp. 201–202).
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — the exists/quality/performs test on the S3⊠S1 accountability channel, and the "poor System 3*" pathology (P2.11).
- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management), and **Hoverstadt** (ch. 7; the pathology bestiary).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- For the set's thesis on AI, System 2 and human agency, see **Designing Freedom, Again** (ch. 9) — the coda.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1981). *Brain of the Firm* (2nd ed.). Chichester: Wiley.

Companion Tautai chapters: Chapter 3 (The Structured Network Organization); Chapter 15 (Spanning the Adaptive Space).

Chapter 13

The Republic of Teams

Ending the War Between Center and Edge with VSM Recursion

In Brief

As organizations scale their intelligence networks, one conflict grows with them: the fight between corporate center and local units over who owns strategy, sensing, and identity. The center sees local strategizing as fragmentation and duplicated effort. The units see central strategy as blindness to their markets. Most organizations resolve this fight by force — centralize, then decentralize, then centralize again — and each swing destroys the capabilities the previous swing built.

The Viable System Model dissolves the fight with its most elegant idea: recursion. Every viable unit contains a complete set of the five systems — its own operations, coordination, optimization, sensing, and identity — nested inside the larger whole like a vaka inside a fleet. The question is never *whether* a local unit should have its own strategy function; if the unit must survive contact with its own market, it already needs one. The real design questions are which decisions live at which level, how the levels talk to each other, and when the higher level may intervene. This showcase turns the center-versus-edge war into that design conversation — and shows why the answer is the structural foundation of Distributed Response.

Two Letters from the Same Company

The two positions in this war state themselves best in the first person — distilled here into the letters their holders might write.

From a divisional managing director to a friend: *"Group strategy has decided our division's priorities for the next three years. The people who wrote them have visited our market twice. We employ four hundred people who talk to our customers every day, and our role in strategy is to receive it."*

From a group strategy director, same company, same month: *"The divisions each run their own strategy exercises, their own market research, even their own innovation labs. We pay for the same work five times, the results contradict each other, and when we try to align them, we're accused of ivory-tower interference. Nobody is steering the whole."*

Both letters are accurate. Both describe real waste and real blindness. And both writers share an assumption so deep neither can see it: that strategy is a single thing that must live in a single place — so the only question is *where*, and the answer must be won.

This assumption produces the oldest pendulum in corporate life. Centralize: gain coherence, lose market contact. Decentralize: regain contact, lose coherence. Each swing takes two to three years, burns a reorganization's worth of energy and trust, and dismantles what the previous swing built. Organizations that swing on a decade rhythm never accumulate either capability.

The Viable System Model's answer to "should strategy live at the center or the edge?" is that the question is malformed — like asking whether a fleet or its vessels should have hulls.

Recursion: Viability All the Way Down

Chapter 1 introduced the five systems every viable organization needs: value-creating operations (System 1), coordination (System 2), operational optimization (System 3), external sensing and adaptation (System 4), and identity and policy (System 5). What this showcase puts to work is the model's structural masterstroke, Beer's recursion principle:

Every System 1 unit of a viable system must itself be a viable system.

Each operational unit — a division, a country organization, a Dynamic Team — is not a fragment that borrows its brain from headquarters. If it faces an environment of its own — its own customers, competitors, regulators — then it needs all five systems of its own, sized to its scale: its own operations, its own coordination, its own optimization, its own sensing of *its* environment, its own working identity. Nested inside the next level up. Like a vaka that is a complete vessel — hull, sail, navigator, crew — and also a member of a fleet that has formations, shared destinations, and a fleet identity of its own.

Seen through recursion, the two letter-writers stop contradicting each other:

- The divisional director is right that his division needs its own System 4 — nobody at group level can sense his market at the resolution his decisions require. Recursion says: he does not need to *win* this against the center. His S4 is not a rebellious duplicate of group strategy; it is a structural necessity at his level.
- The strategy director is right that five uncoordinated strategy exercises are waste — but the cure is not confiscating them. It is connecting them: the group-level System 4's job is not to replace divisional sensing but to sense what *only the whole* can see — patterns across divisions, threats to the portfolio, opportunities that fall between unit boundaries — and to use the divisional S4s as its richest sensor array.

The war over ownership becomes a design conversation about *levels*. That conversation has three parts.

Design Conversation 1: What Belongs at Which Level?

The recursion principle comes with a discipline, Beer's autonomy principle: the higher level absorbs only what the lower level cannot handle — and intervenes in a lower level's affairs only when the viability of the whole is at stake. Not when the center disagrees. Not when the center could do it better. Only when the whole is threatened.

This gives the center-edge negotiation an objective test to replace the power struggle. For any contested decision domain — pricing, hiring, product scope, brand, technology platform — ask:

1. **Does this decision require information that only exists locally?** Then it belongs at the unit level, full stop. Central pricing committees fail not because central people are less smart but because price-relevant variety lives at the edge.
2. **Does this decision create effects that cross unit boundaries?** Shared platforms, shared brand, shared risk pools — these are candidates for the higher recursion level, or for System 2 coordination between units without any escalation at all.
3. **Does getting it wrong locally threaten the viability of the whole?** Compliance exposure, existential financial risk, identity-defining commitments — the higher level holds these, and should say so explicitly rather than policing them ad hoc.

Most center-edge fights, run through this test, dissolve into embarrassingly clear answers. The fights persist only while the test is missing — because then every decision is contested on the general principle of *who rules*, which no one can safely lose.

Note what this test does to the meaning of intervention. When the higher level does step in, it can name which criterion it is invoking — and the unit can challenge it on that ground. Intervention becomes reviewable instead of arbitrary. That single change removes most of the resentment that fuels the decentralization backlash.

Design Conversation 2: How Do the Levels Talk?

Nesting without dialogue is just distributed loneliness. Recursion works when specific channels connect the levels — and each channel pairs functions of the *same type* across levels.

S4-to-S4: The Sensing Federation. The unit-level sensing functions and the group-level System 4 form a standing network — the intelligence network of *The Tautai Principle* (ch. 9) given structural legitimacy. Unit S4s feed local weak signals upward *unaggregated enough to keep their texture* (the express-lane logic of the Gatekeeper's Dilemma, applied between levels).

Group S4 feeds back cross-unit patterns no unit could see alone: "three of you are reporting the same customer behavior shift — this is bigger than any of your markets." This is Distributed Intelligence operating across recursion levels, and it is the most direct answer to the strategy director's duplication complaint: the exercises stop being duplicates when they are federated.

S5-to-S5: The Identity Cascade. *The Tautai Principle* (ch. 10) argued that identity is emergent from consistent action and cannot be copied. Recursion adds: identity is also *nested*. A division has a working identity of its own — and it must be a legitimate local expression of the group identity, not a contradiction of it. The S5-to-S5 conversation sets the non-negotiable identity core that holds everywhere, and explicitly frees everything outside that core for local expression. Units that know exactly which few things are sacred stop fearing that any local deviation will be punished — and centers that have named their sacred few stop fighting battles over the profane many.

S3-to-S3: The Resource Bargain, Stacked. The negotiation format from the previous showcase repeats at every level: group S3 strikes bargains with divisions; divisional S3s strike bargains with their teams. Explicit discretion spaces, minimal reporting, sporadic audit — the same contract, recursively applied. When bargains are stacked cleanly, accountability aggregates without surveillance aggregating.

Design Conversation 3: When the Levels Disagree

Even well-designed levels will conflict — the group S4 sees a portfolio move that damages one division; a division's identity drifts from the group's. Recursion does not eliminate these conflicts. It makes them *addressable*, because each one now has a named type and a home venue:

- A **subsidiarity dispute** ("this decision belongs at my level") goes to the three-question test above, argued on information, cross-boundary effects, and whole-system risk.
- A **sensing dispute** ("your pattern reading contradicts our local evidence") goes to the S4 federation — and is treated as a signal in itself, usually the most valuable kind. When the levels disagree about what the environment is saying, someone is about to learn something.
- An **identity dispute** goes to the S5 cascade, where the question is never "who wins" but "is the identity core mis-specified, or is this genuinely outside it?"

Compare this with the undesigned alternative, where every disagreement escalates as a generic power conflict to whoever sits highest — the schizophrenia pathology of *The Tautai Principle* (ch. 5), with its fragmented communication and incoherent goals, now running *between* levels instead of within one.

Why This Is the Foundation of Distributed Response

Strategic adaptivity requires the organization to respond at the speed of the unit closest to the signal — Distributed Response. But distributed response without recursion is a contradiction: a unit cannot respond to what it senses unless it has its own sensing (S4), its own discretion (bargained S1–S3), and enough identity (S5) to respond in a way the whole can live with. Every capability *The Tautai Principle* builds — Scanning Systems, Dynamic Teams, intelligence networks, narrative steering — reaches its scale limit at exactly the point where the organization refuses to let its units be viable.

The payoff is a structure that ends the pendulum. Centralization and decentralization stop being opposed eras and become simultaneous properties of different decision domains: coherent where the whole demands it, autonomous everywhere else, with the boundary between the two explicit, tested, and renegotiable. What emerges is less an empire ruled from the center than a republic of teams: self-governing units under a shared constitution, with the few federal matters named and everything else left to the provinces. The organization stops swinging and starts sailing in formation.

Seen in Practice: One Retailer, Five Market-Research Bills

An international grocery retailer ran the old pendulum. The group had centralized buying to capture scale; the German and Spanish country organizations complained that group ranges ignored local taste and quietly ran their own assortment teams — duplicated work the center called fragmentation.

The recursion reframe treated each country organization as a viable system in its own right — its own operations, coordination, sensing and identity — nested inside the group. The design conversations then settled which decisions live where: private-label sourcing, payment terms and global supplier standards at group, where scale and non-negotiables belong; local range, promotions and store format at country, where market contact lives; and a written rule for the rare collision — group sets the frame, countries fill it. A twice-yearly assortment negotiation replaced the annual centralize-or-decentralize argument. The group stopped paying five times for the same market research, because each country's sensing function now fed a shared group-level one instead of hiding from it. Neither of the two letter-writers had been wrong. The question "should assortment live at the center or the edge?" had simply been malformed, and nested viability answered it level by level.

Which method finds this. Beer's trio discipline catches this dysfunction at the modeling stage: never model a level without the one above and the one below it, so neither letter-writer's view can pose as the whole (ch. 1). Pérez Ríos would unfold the organization along several criteria at once — geography, product, market — and check the result against his structural pathologies of vertical unfolding, several of which name exactly this case

of missing or ill-defined recursion levels (ch. 4). Hoverstadt reads the organization fractally — the same design rules at every scale — so a center that denies its units their own sensing fails a test it must also pass itself (ch. 7). Here nobody holds the pen alone: each level writes its own five systems, and the level above countersigns only the boundary between them.

Monday Morning Actions

- **Map one unit's five systems.** Pick one division or major team. Ask: where is its S4? Its working S5? If the answers are "headquarters does that," you have found viability gaps — places where the unit senses nothing and decides nothing, and therefore cannot respond to anything. Owner: unit lead with Head of Strategy.
- **Run the three-question test on your sorest fight.** Take the most contested center-versus-unit decision domain of the past year. Apply the three questions in a joint session with both parties. Write down the answer and who intervenes on what grounds. Owner: CEO or COO.
- **Federate two S4s.** Connect group strategy with the sensing function of one unit in a monthly signal exchange — unaggregated local signals up, cross-unit patterns down. Measure whether either side learns something it did not know. Owner: Head of Strategy.
- **Name the identity core.** Have the executive team write the list of things that must hold at every level of the organization — and cap it at five. Publish it with the explicit statement that everything else is open to local expression. Owner: CEO.

The apprentice asked: "Master, when the fleet sails together, why does every vaka still carry its own navigator? Would not one navigator for the fleet be enough?" The master laughed. "One navigator can read one horizon. The fleet spans many. Each vaka reads its own water, chooses its own line through the swells — and holds the formation because every navigator carries the same stars in memory. A fleet with one navigator is one storm away from being no fleet at all."

Sources & Related Chapters

Recursion, the System-in-Focus and nested viability used here are diagnosed and designed with the methods profiled in the VSM method-chapter set:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — recursion as an exact structural relation, the System-in-Focus, and the container-system-contents trio.
- Pfiffner, *The Neurology of Business* (ch. 2) — recursion levels R+1 / R0 / R–1, and "Overlooking Nesting" as a cardinal error of organizing.
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — recursive analysis and "unfolding the complexity" of nested primary activities.
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — multi-criteria recursion, the Recursion Levels–Critical Factors Matrix, and the structural pathologies of vertical unfolding (P1.1–P1.4).
- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management), and **Hoverstadt** (ch. 7; the pathology bestiary).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- For the set's thesis on AI, System 2 and human agency, see **Designing Freedom, Again** (ch. 9) — the coda.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1981). *Brain of the Firm* (2nd ed.). Chichester: Wiley.

Companion Tautai chapters: Chapter 5 (The Organizational Nervous System); Chapter 9 (Scaling to Intelligence Networks); Chapter 15 (Spanning the Adaptive Space).

Chapter 14

Steering Without Shouting

Why More Communication Produces Less Alignment — and What System 5 Is Actually For

In Brief

Leadership teams broadcast more than ever — town halls, cascades, newsletters, OKR reviews — while alignment measurably declines: the communication overload paradox. The usual diagnosis is a messaging problem: wrong format, wrong frequency, wrong words. This showcase offers a structural diagnosis instead. In Viable System Model terms, the paradox arises when System 5 — identity, purpose, and policy — tries to align the organization through *instruction*, a channel that can never carry enough variety to cover thousands of daily local decisions. The alternative channel is *identity transmitted through narrative*: shared stories that let any team resolve a local trade-off the way the organization would, without asking.

This reframing turns a perennial soft topic — storytelling, purpose, culture — into hard cybernetic engineering with a measurable output: the percentage of local decisions that align with strategic intent without escalation. It also completes the Adaptive Space — the arena where operations and strategy contend, developed in *The Tautai Principle* (ch. 15): System 5 can only moderate the System 3–4 conflict at the speed adaptation requires if its verdicts are *anticipatable* — if both sides already know the story well enough to predict how the organization would decide.

The CEO Who Communicated Everything Except Direction

The CEO of a mid-sized technology firm — call her Marta — was, by every conventional standard, an outstanding communicator. Monthly town halls with open Q&A. A weekly video message. Quarterly strategy cascades, professionally produced, translated into three languages. An internal channel where she personally answered employee questions.

Then her leadership team ran a simple alignment probe: they took twenty recent operational decisions from around the company and asked, for each, *would we have decided the same way?* The result: eleven out of

twenty. Barely better than a coin flip — after a year of the most intensive communication program in the company's history.

Marta's first instinct was to communicate *more*. Her head of strategy talked her out of it with one observation: "The teams can all recite the strategy. They just can't *use* it. When the new customer asks for the discount, when the release date collides with the quality bar, when the partner wants an exception — the strategy deck doesn't answer. So they guess, or they escalate."

That distinction — between a message that can be recited and a story that can be used — is the entire subject of this showcase. And the reason it matters is not communication theory. It is channel capacity.

The Arithmetic That Breaks Instruction

Return to Ashby's Law — only variety can absorb variety (ch. 1) — and apply it to the top of the organization.

An organization of two thousand people makes tens of thousands of decisions every week that touch strategic intent: pricing exceptions, quality trade-offs, hiring calls, customer promises, technical shortcuts taken or refused. Call it the *decision variety* of the organization. For the organization to act coherently, something must align those decisions with the whole.

The instruction channel — policies, approvals, cascaded objectives, escalations — tries to align them by *specifying* them. But channel capacity is brutally finite: leadership attention, policy pages, meeting hours. Confronted with decision variety in the tens of thousands, the instruction channel has three failure modes, and every organization exhibits at least one:

- **Specify more** — policy proliferation, until the manual answers everything and no one can find anything.
- **Escalate more** — the queue outside the executive office, where decisions wait for the only people entitled to make them, at hierarchy speed in a market-speed world.
- **Broadcast more** — Marta's path: flood the organization with context and hope alignment emerges. It does not, because broadcast increases the *volume* of the channel, not its *usability at the point of decision*. A team facing a specific trade-off cannot query a town hall.

This is the communication overload paradox, restated structurally: **alignment demand grows with decision variety; instruction capacity does not.** More communication of the same kind cannot close the gap. The channel is the wrong shape.

What System 5 Actually Transmits

The Viable System Model names a fifth system — identity, purpose, policy — whose day job includes moderating the System 3–4 conflict. It is tempting

to picture System 5 as the top of the pyramid: the board, the executive committee, the people who decide the biggest things. That picture is wrong in an instructive way.

System 5's real function is *closure*: it is the mechanism by which the organization remains one organization while everything else adapts. And its most powerful mode of operation is not deciding cases. It is making decisions *unnecessary* — by transmitting the identity that lets any unit, at any recursion level, resolve its own cases the way the whole would.

Look at the phenomenon inside every organization that has genuine identity: the veteran team lead who, facing a novel situation, says "that's just not how we do things — we'd rather lose the deal than promise what we can't ship," and is *right*, and everyone knows she is right, though no policy covers the case. That sentence is System 5 operating at the frontline, at zero channel cost, at market speed. Identity answered a question that instruction never anticipated.

This is variety attenuation of the most elegant kind: instead of matching tens of thousands of decisions with tens of thousands of instructions, the organization installs *one decision pattern* in every head, and the pattern generates the answers locally. *The Tautai Principle* (ch. 10) argues that identity is the one thing competitors cannot copy. Add now: it is also the only alignment channel whose capacity *scales with the organization* — because every member who carries the story is additional bandwidth.

Why Stories, Specifically

If identity is the channel, narrative is the encoding — and the choice of encoding is not aesthetic. Stories have three properties that mission statements, value lists, and strategy decks lack.

Stories carry trade-offs. A value list says "quality" and "customer focus" and "speed" — and goes silent exactly when those collide, which is the only moment guidance is needed. A story is *about* the collision: the time we refunded the flagship customer and missed the quarter; the time we shipped late rather than ship fragile; the time we killed our most profitable product because it no longer fit who we were becoming. Stories encode the *ranking*, not just the values. That is why they can be used at the point of decision and a poster cannot.

Stories are generative. A policy covers the cases its author foresaw. A story projects onto cases no one foresaw — the listener maps the new situation onto the old pattern and reads off the answer. This is precisely the property Ashby demands: a compact message that unfolds into high variety at the receiving end.

Stories replicate without the center. Instructions travel down the hierarchy and stop when the cascade stops. Stories are retold peer-to-peer, at onboarding, at lunch, after the crisis — through exactly the informal networks that carry most of what an organization actually knows (mapped in *The Tautai Principle*, ch. 9). The center seeds; the network propagates.

Mavens carry them, connectors spread them, and every retelling is a free repetition the town hall budget never has to pay for.

One warning before the practice: narrative steering without narrative *honesty* is worse than silence. The organization also tells itself counter-stories — about the executive who violated the espoused value and was promoted, about the quality speech followed by the shipped shortcut. Stories are falsifiable by observed behavior, and the counter-story always wins because it carries better evidence. System 5's narrative authority is exactly as strong as its behavioral consistency. This is the identity discipline of *The Tautai Principle* (ch. 10) restated as a communication constraint: you cannot narrate what you do not enact.

Narrative Steering as an Operating System

Turning this from insight into practice means running narrative as an engineered loop — the helm described in *The Tautai Principle* (ch. 12), now given its cybernetic job description.

Curate decision stories, not success stories. The steering canon is not the highlight reel. It is a small, deliberately maintained set of stories — a dozen is plenty — each preserving a hard trade-off and its resolution. For each strategic priority, ask: *what story teaches a team how to decide when this priority collides with another?* If no story exists yet, the priority has not yet been lived — a finding in itself.

Retire stories that no longer steer. Identity evolves, and a story canon that only grows becomes a museum. When the strategy genuinely shifts, the old story must be explicitly retired and replaced — ideally by narrating the shift itself: "we used to be the company that never said no to a customer; here is the moment we learned to."

Harvest from the edge. The best steering stories are not written by communications; they are *found* — in the express-lane signals of the Gatekeeper's Dilemma, in System 3–4 collaborations, in incident reviews. When a team resolves a novel trade-off well, System 5's job is to notice, canonize, and retell. The organization steers itself with its own best moments.

Close the loop with the alignment probe. Narrative steering has a hard metric: the share of sampled local decisions consistent with strategic intent *without escalation* — Strategic Alignment, as *The Tautai Principle* (ch. 5) names it. Run Marta's twenty-decision probe quarterly. When a cluster of misaligned decisions appears, do not write a policy — find the missing or broken story. When a decision was escalated that the story should have answered, ask why the story did not reach that team, and trace which informal network failed to carry it.

Completing the Adaptive Space

This showcase began as a communication problem and ends as the missing piece of the Adaptive Space. There, System 5 moderates the productive

conflict between operational management (System 3) and strategic intelligence (System 4): framing tensions, arbitrating resources, facilitating exchange. But a System 5 that must convene to decide every case becomes the bottleneck of its own Adaptive Space — moderation at committee speed, adaptation at market speed.

Narrative closes the gap. When the identity story is vivid enough, System 3 and System 4 can *anticipate* the moderation: both sides know the story well enough to predict how the organization would weigh this quarter's efficiency against next year's option. Most conflicts settle at the table, by reference to the shared narrative, without ascending. System 5 then reserves its live attention for the cases that genuinely rewrite the story — which is exactly the learning-integration role the Adaptive Space assigns it.

An organization steered this way exhibits what *The Tautai Principle* (ch. 2) calls dynamic stability: everything operational adapts continuously, while identity — the thing that makes the adaptations *coherent* — changes slowly and deliberately. The helm is quiet. The vessel still answers it.

Seen in Practice: The Front Desk That Didn't Need the Manual

A mid-market hotel group had a forty-page service manual and still could not get consistency. A front-desk clerk facing an angry guest at midnight either escalated (slow) or improvised (off-brand), and more rules made it worse — nobody could find the relevant page at 11pm.

The shift was from instruction to a story people could use — a story of the kind every service company polishes, and worth retelling anyway: a single, retold account of a night manager who found a family stranded by a cancelled flight, comped a room and a meal without calling anyone, and was backed the next morning — because "we are the place that solves the traveller's worst night." That story did what the manual could not: it let a clerk resolve a novel cost-versus-guest trade-off the way the company would, without a lookup and without an escalation. Leadership stopped adding policies and began curating a small canon of such stories, each carrying one identity-defining trade-off. Alignment — the share of frontline judgment calls that matched intent without escalation — roughly doubled over three quarters. Not from more communication, but from stories staff could actually use at the point of decision.

Which method finds this. Espejo and Reyes would start where this showcase ends: for them, naming the organization — settling what it is, for whom, in whose eyes — is the first and most consequential diagnostic act, and TASCOI exists to force that naming into the open (ch. 5). Schwaninger gives the same concern its own altitude: normative management, the third and slowest clock, which legitimizes what the faster clocks of strategy and operations are allowed to do (ch. 6). And Beer supplies both the seat — System Five as the closure that keeps one organization one — and the test: the purpose of a system is what it does, so the story leadership tells is falsified daily by the decisions the organization actually makes (ch. 1). In

this scenario the pen sits with whoever curates the story canon — and the probe exists to check that the told story and the enacted one are still the same book.

Monday Morning Actions

- **Run the twenty-decision probe.** Collect twenty recent operational decisions; have the leadership team assess: would we have decided the same way? Your score is your System 5 channel health. Below fifteen, you have a steering gap no additional broadcasting will close. Note the circularity: the team grading the alignment is the team whose story is under test — blunt it by having the units' own people score the same twenty decisions and comparing the two readings. Owner: CEO.
- **Write the three collision stories.** For your top three strategic priorities, find and write the story that teaches how each priority wins or loses against another. Real events, real trade-offs, short enough to retell. If a story is missing, note which priority has never been tested — and treat that as a sensing result. Owner: Executive team, one story each.
- **Find one counter-story and fix its source.** Ask five trusted people at the front line: "What story do people tell here that contradicts the official one?" Then address the *behavior* the counter-story reports — not the story. Owner: CEO, without delegation.
- **Canonize one edge story.** Find a team that recently resolved a hard trade-off well, and have its story told — by the team, not about them — in the next all-hands. You are seeding the network, not producing content. Owner: Head of Communications with Head of Strategy.

The apprentice asked: "Master, on the long crossing you sleep through the night watches. How does the crew hold the course while you dream?" The master navigator smiled. "Before every voyage, I tell them of the crossing where we followed the false current, and what it cost us, and how we found the true one again. I do not leave them a list of orders — no list survives first contact with the sea. I leave them the story. When the water changes in the dark, they do not wake me. They ask the story what I would do — and it answers."

Sources & Related Chapters

System 5 — identity, closure, and the narrative that aligns local decisions — is diagnosed and designed with the methods profiled in the VSM method-chapter set:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — System 5 as logical closure and self-reference (the assertion of identity), and the algedonic alert that reaches it.
- Pfiffner, *The Neurology of Business* (ch. 2) — System 5 as identity and norms, and the S5–S3/4 interaction (Ch. 13.5).
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — System 5 as the shared ethos that lets units decide the way the whole would.
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — System 5 identity pathologies — ill-defined identity (P2.1) and institutional schizophrenia (P2.2).
- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management), and **Hoverstadt** (ch. 7; the pathology bestiary).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- For the set's thesis on AI, System 2 and human agency, see **Designing Freedom, Again** (ch. 9) — the coda.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1981). *Brain of the Firm* (2nd ed.). Chichester: Wiley.

Companion Tautai chapters: Chapter 5 (The Organizational Nervous System); Chapter 10 (Identity as Competitive Advantage); Chapter 12 (Steering with Stories); Chapter 15 (Spanning the Adaptive Space).

Chapter 15

Where the Machine Belongs

Placing AI Across the Five Systems Without Automating the Judgment That Makes You You

In Brief

The other showcases in this part each take a chronic structural tension and turn it into a designed exchange. This one deliberately breaks that pattern, and says so, because the fight it addresses is not a tension in any one seam at all. The argument over artificial intelligence — automate everything, or protect human judgment everywhere — is an *allocation* question across all five systems at once: in which functions can AI absorb variety, and in which is human judgment constitutive?

Framed as a totality — humans versus machines, IT versus the business — every AI initiative becomes a referendum on trust, jobs, or ideology, and the actual placement decisions get made by procurement momentum and vendor demos instead of by design. The Viable System Model dissolves the referendum by giving the question a grid. AI absorbs variety well in System 1 execution, System 2 coordination, System 3 monitoring, and System 4 pattern detection. It does not hold purpose: interpreting what a pattern *means for us* (inside System 4) and committing to the trade-offs that define identity (System 5) are the acts that make an organization this organization, and handing them over is not automation — it is abdication.

This showcase shows how to convert the referendum into a maintainable design. The instrument is a written **judgment boundary** for each placement: what AI absorbs, what humans decide, and how the handover between them is engineered — reviewed on a rhythm, because the boundary moves as the technology moves. It is the practical companion to the coda's argument: the machine can be a liberty machine or its opposite, and the difference is drawn, function by function, by whoever holds the pen.

The Referendum That Wasn't a Vote

Watch how the argument usually arrives. The board issues a mandate — "AI in every process by year end." The works council hears the same sentence translated — "judgment removed from every process" — and digs in. Both then behave accordingly, and both are responding to the same category

error: they are treating a question about *where* as though it were a question about *whether*.

One camp pushes automation everywhere, because it cannot afford not to. The other defends human involvement everywhere, because judgment cannot be delegated. Each argues in totalities, so every proposal becomes a proxy war over trust and jobs. And while the two camps fight the referendum, the placement decisions get made anyway — not by either of them, but by whichever vendor demo landed best. The org chart of AI usage ends up recapitulating the sales funnel rather than the organization's viability requirements.

You cannot resolve this by winning the referendum, because there is no single thing to vote on. Automating document extraction in the claims queue and automating the "what should we become" conversation in the boardroom are not the same decision at different volumes. They are different decisions about different functions, and the only way to make them well is to stop arguing about AI in general and start asking, of each specific placement, a much narrower question.

Why the Question Is Malformed

"Should we automate this?" is malformed in the same way "should we delegate?" is malformed. Delegate *what*, to *whom*, with what authority to object? The useful question is always about a specific function and a specific handover.

Underneath sits the same cybernetic fact that governs every other showcase in this part: Ashby's Law of Requisite Variety. The environment throws more variety at an organization than any part of it can absorb, so the work of design is deciding *what absorbs what*. AI is, at bottom, a new and very large capacity to absorb certain kinds of variety — patterns, transactions, documents, correlations — at a scale no human can match. That makes it a variety engineer of extraordinary reach. It also means the right question is never "automate or not," but: *in which system would this AI sit, and what variety would it absorb there?*

Ask it that way and the humans-versus-machines frame falls apart, because the answer is different in every system. The machine belongs in some functions and is corrosive in others, and the difference is not about how good the model is. It is about what each function is *for*.

The Allocation Grid

Walk the five systems and the pattern is clear.

System 1 — execution. The value-creating work carries massive routine variety: orders, tickets, transactions, documents. This is pattern-work, and AI absorbs it well — extraction, routine settlement, first drafts. What stays human is the exceptions, and the client-facing moments where the interaction *is* the product.

System 2 — coordination. Scheduling, deconfliction, translating one unit's standards into another's. Rule- and pattern-work again, and AI absorbs it easily. This is the least glamorous placement and the highest-yield one, because System 2 is chronically underbuilt — the empty coordination seam that organizations usually leave to heroics. It is also, as the coda argues, the placement with the deepest structural consequence: a strong System 2 lets autonomous units coordinate without a boss, which is what makes flatter structures hold.

System 3 — monitoring. Aggregation, anomaly flagging, audit preparation, forecasting inputs. AI absorbs the *monitoring* well. But the resource bargain itself — the negotiation between accountable parties over what autonomy is granted for what accountability (Showcase 3) — is a human act, constitutively, because it is a commitment between people who will answer for it.

System 4 — intelligence, and here the boundary runs straight through the middle of the system. System 4 does two different jobs. *Pattern detection* — scanning volumes no human could read, surfacing anomalies and correlations — is exactly what AI is best at, and it extends the Scanning System of *The Tautai Principle* (ch. 4). But *interpretation* — deciding what a pattern *means* for this organization, which weak signal matters, what it implies for strategy and identity — is constitutively human, because meaning is assigned relative to purpose, and purpose is not in the training data.

System 5 — identity. The trade-offs between now and next, between values in genuine tension, the "what we would never do" calls. These are constitutively human. An organization that delegates them has not automated System 5; it has dissolved it. AI's only legitimate role here is adversarial — stress-testing a proposed trade-off, never making it.

The word doing the work is *constitutive,* and it has to mean something more durable than "AI isn't good enough yet," or the whole line erodes with the next model release. It means this: in some functions, the *act of deciding is itself what makes the organization this organization*. Assigning meaning and committing to trade-offs are identity-performing acts. A better model changes how much *support* those acts can draw on; it never changes *whose* acts they are. Content becomes cheap, as *The Tautai Principle* (ch. 11) argues, and judgment becomes the scarce, precious thing — a scarcity you protect structurally, not with slogans. Delegating the constitutive functions is not efficiency. It is abdication.

The Judgment Boundary

Naming where AI belongs is diagnosis. The design object is the **judgment boundary**: for each placement, a written statement of what AI absorbs, what humans decide, and — the interesting part — how the handover between them is engineered.

The seam is where the real work is, because a boundary without a designed handover degrades into rubber-stamping within a quarter: automation with an alibi signature. A boundary worth the name specifies four things. What the model hands over, with its confidence levels *and its discard pile* — the "doesn't fit" cases it could not classify. What the human can override, and whether overriding is cheap enough that it actually happens rather than being nominally available and practically impossible. What triggers escalation. And who reviews the boundary, and when — because as the technology's absorption capacity grows, the boundary moves, and a boundary without a review date is simply a future referendum you have agreed to hold by accident.

One item on that list is load-bearing across every placement: the discard pile. Every AI placement optimizes against the patterns of the past, which is fine in Systems 1 through 3 as long as the anomalies keep flowing upward — because the model's "this doesn't fit" pile is not noise to suppress, it is System 4's feedstock. The quiet catastrophe is the placement that silences it: the triage model that files every edge case into an existing category, so the weak-signal stream dries up and nobody decided that it should. Placed unguarded in System 4 or 5 territory, the same optimization is fatal in a different way — the organization gets efficiently, confidently better at yesterday. The boundary's escalation triggers and its mandatory discard-pile route are the anti-entrenchment mechanism. They are how you keep the machine from quietly steering while everyone believes a human still is.

There is a second test a boundary must pass, and it is not an effectiveness test. A placement can clear every hurdle above — clean handover, cheap override, live discard pile — and still shift who decides, who sees, and who is seen. That is the coda's criterion (ch. 9): the question is never only whether the machine works, but whether it enlarges human agency or concentrates it. Which is why the works council from the opening scene belongs in the boundary-writing, not just in the referendum it was handed. Judging placements by signal streams and cycle time alone is precisely the technocratic move the coda warns against — variety engineering with the politics left out.

Seen in Practice: The Claims Department That Fed Its Own Strategy

An insurer's claims unit deployed AI three times, and the three placements are the whole argument in miniature.

The first was in System 1: document extraction and routine settlement below a threshold, with humans handling exceptions and hardship cases. Cycle time fell by around 60% (the figure is the vendor's), and complaints fell too, because the human minutes now went where they mattered. A clean absorption.

The second was in System 3, and it came with a guard: anomaly detection across claims patterns for audit support, with the model's monthly "strange but not fraud" pile routed, unfiltered, to the strategy team. That

pile is where they first noticed a new mobility-behaviour pattern that later reshaped a product line — System 4 fed by System 3's discards, exactly as designed. Without the guard, that signal dies in the model's reject folder.

The third crossed the line. Leadership piloted an AI-generated "strategic implications" report built from market data, and for two quarters the strategy meetings quietly reorganized themselves around discussing the model's takes. The room anchored on machine-plausible consensus; weak signals got argued *away* because the model had scored them low-relevance. This is interpretation automated — the constitutive act of System 4 handed to a pattern-matcher. The correction was not to switch the tool off but to demote it: the model stayed as a pattern-supplier, its detection lists intact, while interpretation returned to a human rhythm — with the model's discard pile explicitly on the agenda, and the whole boundary written down with review dates.

Same technology, three placements, three outcomes. Move the story to a newsroom and the seams shift but the lesson holds: auto-transcription in System 1 frees reporters for investigation; an AI scheduler quietly fills the empty System 2 coordination seam; and an engagement-trained "what should we cover" ranker, left to set the front page, drifts the outlet's identity toward whatever maximizes clicks — System 5 dissolving — until it is demoted back to an input and the editorial line returns to humans. The danger simply migrates to whichever system holds the constitutive act.

Failure Modes to Design Against

Judgment automated. The model drifts from detection into interpretation, or from stress-testing into deciding. Watch the boundary between pattern and meaning; it is the one that matters, and it is invisible until you have named it.

Absorption refused. Toil defended in the language of judgment — analysts spending days consolidating reports "because leadership wants a human touch," while the actual judgment work gets thirty minutes at the end. Control theatre is as much a design failure as reckless automation; it just fails quietly, in System 2 and System 3.

The silenced anomaly stream. A placement that files every "doesn't fit" into an existing bucket has cut System 4's feedstock. The discard-pile route is not optional plumbing; it is the load-bearing wall.

The rubber-stamp. A handover where a human nominally reviews the model's output but override is slower than compliance is not a boundary — it is automation with a signature line. Test it: is overriding genuinely cheaper than going along? If not, redesign with sampling, confidence-gated routing, or batch limits.

Which method finds this. The capstone's allocation logic is built for exactly this question — matching capacity to function, system by system, before any tool is bought (ch. 8). Beer's variety laws set what the machine

may absorb: where variety is routine and countable, absorption is legitimate; where it is meaning, it is not (ch. 1). Espinosa and Walker add the warning this showcase's works council embodies — participative legitimacy is slow, and it is precisely what AI makes cheapest to skip (ch. 3). And the coda supplies the criterion every placement must finally answer to: does it enlarge human agency or concentrate it (ch. 9)? Here the pen belongs to whoever writes the judgment boundaries — and the design question is whether that is a few hands or many.

Monday Morning Actions

- **Reframe one referendum.** Take one AI initiative that has become a for-or-against fight and re-ask it as a placement question: which system would this sit in, and what variety would it absorb there? Owner: whoever sponsors the initiative.
- **Map the portfolio onto the five systems.** Inventory current and proposed AI deployments, place each in its system, and tag it: sound, judgment-automated, absorption-refused, or seam-undesigned. Then run the reverse pass — for each system, what variety is nobody absorbing? (System 2 is reliably empty of AI and full of toil.) Owner: COO with the AI/data lead.
- **Draft one judgment boundary.** For the highest-stakes placement, write it down: what AI absorbs, what humans decide (named as decisions, not "oversight"), the handover design including the discard-pile route, and a review date with an owner. Owner: the deployment's actual owner.
- **Name the anomaly stream.** For every placement in Systems 1–3, state which friction signal must keep flowing upward, and who notices if it stops. Owner: Head of Strategy (System 4).

The apprentice asked: "Master, this new instrument counts the stars faster than any of us, and never tires. Should we not let it choose our heading?" The master navigator turned it over in her hands. "Let it count the stars — it counts them better than I do, and my eyes are tired. Let it call the swells, and the birds, and the drift. But the heading is not a counting. The heading is what we are, made into a direction: where this crew has sworn to go, and what we will not do to get there. Give the instrument the counting and it makes us swifter. Give it the heading and there is no longer an us for it to carry — only a very fast canoe, going wherever the last star pointed."

Sources & Related Chapters

The variety-allocation logic across Systems 1–5 rests on the primary VSM literature; the method chapters in Parts I and II show how each placement is diagnosed and chosen in practice:

- Beer, *Diagnosing the System for Organizations* (ch. 1) — variety, requisite variety (Ashby's Law), and the five systems' functions, in the originator's own worksheets.
- Pfiffner, *The Neurology of Business* (ch. 2) — measuring complexity as variety, and distributing mission-critical tasks across the control functions (Steps III–V).
- Espinosa & Walker, *A Complexity Approach to Sustainability* (ch. 3) — variety engineering of the S1–environment loop.
- Pérez Ríos, *Design and Diagnosis for Sustainable Organizations* (ch. 4) — the exists/quality/performs triple test applied per component, and the functional-pathology catalogue organized system by system.
- The set is completed by **Espejo & Reyes** (ch. 5; VIPLAN / TASCOI), **Schwaninger** (ch. 6; the three logical levels of management), and **Hoverstadt** (ch. 7; the pathology bestiary).
- For choosing among these methods for a given problem, see **Choosing a Method** (ch. 8) — the capstone comparison and decision guide across all seven.
- This showcase is the practical companion to **Designing Freedom, Again** (ch. 9) — the coda, where the judgment boundary becomes the design that decides whether AI enlarges human agency or concentrates it.

Primary VSM literature: Beer, S. (1985). *Diagnosing the System for Organizations*. Chichester & New York: Wiley. · Beer, S. (1979). *The Heart of Enterprise*. Chichester: Wiley. · Beer, S. (1974). *Designing Freedom*. Toronto: CBC / Chichester: Wiley.

Companion Tautai chapters: Chapter 4 (Building the Scanning System); Chapter 5 (The Organizational Nervous System); Chapter 11 (AI, the New Landscape, and the Scarcity of Judgement).

Glossary

The Twenty Terms

Beer promised that his own handbook could be worked through with "less than twenty technical terms." This book keeps the same discipline: the twenty below are all it assumes. Definitions follow the usage of this book and its sources; where authors differ, the difference is noted in the chapter concerned.

Viable system — a system capable of independent existence: of maintaining its identity and adapting in a changing environment. "Sustainable," in most of the books profiled here, means exactly this — not ecology (the exception is Espinosa & Walker, ch. 3).

Variety — the number of possible states of a system; Beer's measure of complexity. Everything in the model is an answer to the question of who absorbs how much of it.

Requisite variety (Ashby's Law) — only variety can absorb variety: a regulator must command at least as much variety as the situation it regulates. The law behind every structural argument in this book.

Attenuator — a design element that reduces the variety reaching a point in the system (a report, a filter, an exception threshold, an aggregation).

Amplifier — a design element that increases the variety a point in the system can deploy (a tool, a delegation, a standard, a team).

Homeostat — a pair of parts holding each other in balance through mutual variety exchange; in this book most often the System 3–System 4 balance between today and tomorrow.

Recursion — the nesting of viable systems: every viable system contains viable systems and is contained in one, with the same five-system structure at every level. Hoverstadt's "fractal" names the same fact.

System-in-Focus — the one level of recursion you have chosen to model, always held together with the level above and the level below (Beer's "trio discipline").

System 1 — the operational units that *produce* the organization — the parts that could, in principle, be hived off and survive. Distinguished from support functions, which could not.

System 2 — the anti-oscillation function: coordination that keeps System 1 units from colliding, explicitly *off* the command axis. It damps; it does not command. The hinge of the coda's AI argument.

System 3 — the "inside-and-now": day-to-day management of the whole operational cluster, holder of the resource bargain, owner of synergy.

System 3* — the sporadic, high-variety audit channel: occasional deep dives into operations (audits, task forces) that close the gap routine reporting leaves. The price of autonomy in the resource bargain.

System 4 — the "outside-and-then": the organ of adaptation, watching the environment and the future, in permanent negotiation with System 3.

System 5 — identity and closure: the function that says what the organization is, is not, and is for; referee of the 3–4 balance. Emphatically not "the boss."

Metasystem — Systems 3, 4 and 5 taken together: the management of the operational cluster, at a higher logical order than the operations — "higher authority is not the issue."

Algedonic signal — Beer's pain/pleasure alarm: a signal that bypasses the ordinary channels to reach the metasystem directly when a threshold is crossed. "The cry is 'wake up — danger!' "

Resource bargain — the deal at the heart of the vertical axis: autonomy granted to a System 1 unit in exchange for agreed accountability, verified by 3* audit rather than by continuous surveillance.

Autonomy — in Beer's exact sense, residual variety: the freedom remaining to a unit to act on its own account within the framework set by the whole. The operational form of what *Designing Freedom* calls freedom.

Autopoiesis — self-production: a viable system's maintenance of its own identity through change (after Maturana and Varela). Its pathology — self-production become an end in itself — is the founding entry of every pathology catalogue in this tradition, from Beer's "Establishment" to Hoverstadt's Organizational Cancer.

Complexity drivers — the criteria along which an organization's variety naturally divides — technology, customer/supplier, geography, time — used to unfold it into nested viable systems. The order in which you apply them is a design decision, not an accident.