


How Organizations Understand

Situation Awareness,
Sensemaking, and Learning
in Practice



How Organizations Understand

Krishan Mathis

Tautai Press

Copyright © 2026 Krishan Mathis
All rights reserved.
Tautai Press
Open draft — 2026-10-11

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

Contents

1 Chapter 1 — Understanding comes first	1
2 Chapter 2 — Perceiving: attention and the signals at the edges	25
3 Chapter 3 — Knowing what is going on: situation awareness	43
4 Chapter 4 — When the picture breaks: sense-making	59
5 Chapter 5 — Orientation: keeping the model alive	75
6 Chapter 6 — Choosing: from understanding to commitment	93
7 Chapter 7 — Where understanding is held	111
8 Chapter 8 — When groups understand more, and when less	129
9 Chapter 9 — Structure as cognitive architecture	149
10 Chapter 10 — Whose account counts: leadership and sensegiving	167
11 Chapter 11 — Learning loops: how organizations create knowledge, and lose it	185
12 Chapter 12 — The skill development cycle	205
13 Chapter 13 — Making sense with machines	223

14 Chapter 14 — Designing the understanding organization	241
15 Appendix A — This book and The Tautai Principle	271
16 Appendix B — Glossary	277

Chapter 1

Chapter 1 — Understanding comes first

A team that kept the picture

The *Palau*, a US Navy ship, was entering San Diego harbor at about ten knots when the engine room lost steam pressure. Within a minute the ship had no propulsion. The generators stopped, the radars spun down, and a little later the gyrocompass failed, the instrument on which every bearing the navigation team took depended.¹ A ship of that size does not stop when its engines do. It carries on into a narrow channel with other traffic in it, and somebody has to know where it is, minute by minute, until it can be anchored.

Edwin Hutchins, the cognitive scientist who was on board, calls the ship the *Palau*. What he watched was a navigation team that had to keep fixing the ship's position while computing by hand what the instrument normally delivered. Nobody had planned for this. The team reorganized its own division of labor on the move, who takes which bearing, who corrects it, who plots, and about twenty-five minutes after the casualty the ship lay at anchor at the intended spot. No single person on the bridge, Hutchins stresses, not the captain, not the navigator, not the quartermaster chief, could have kept control of the ship alone.² The picture of where the ship was sat in no

¹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 1 "Welcome aboard", pp. 1–5 (loss of steam, anchoring, no single person could have done it); gyrocompass failure and the hand computation of true bearings, ch. 8, pp. 317–318.

²Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 1 "Welcome aboard", pp. 1–5 (loss of steam, anchoring, no single person could have

one head. A team held it, with its charts, its instruments and a set of practiced routines, and the picture held when a piece of that system failed.

The story has a second part. When Hutchins later analyzed what the team had done, he found that the new way of dividing the work had been discovered “by the organization before it was discovered by any of the participants”, and that afterwards “the solution ... was not saved in the system”.³ The team had learned something under pressure and had no way of keeping it.

I start with the *Palau* because those twenty-five minutes hold most of what this book is about: an organization that had to know its situation while acting in it, that held that knowledge together rather than in a leader’s head, that kept it alive when the situation broke, and that then let what it had learned slip away.

A crew that lost it

On 5 August 1949 lightning started a fire south of Mann Gulch, a side valley of the Missouri River in Montana. Fifteen smokejumpers of the US Forest Service parachuted in that afternoon and joined a fire guard who had been fighting the fire alone. They were trained and equipped for exactly this job. Thirteen of the sixteen men died.⁴

Karl Weick, in 1993, re-read the disaster and made it a standard case in organization studies.⁵ The crew, in his reading, did not die from a lack of information. They had a picture of the situation, a fire of a familiar kind, to be

done it); gyrocompass failure and the hand computation of true bearings, ch. 8, pp. 317–318.

³Hutchins (1995), ch. 8, p. 351.

⁴Mann Gulch fire, 5 August 1949, Helena National Forest, Montana: 15 smokejumpers plus fire guard James O. Harrison; 13 died; foreman Wagner Dodge survived in his escape fire, Sallee and Rumsey through a crevice in the ridge; the order to drop tools. — Weick (1993), pp. 633, 635 (cosmology/minimal organizations; drop-tools); primary account Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.

⁵Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. <https://doi.org/10.2307/2393339> (p. 633 (cosmology episode; minimal organizations), p. 635 (drop tools as loss of identity)).

approached and contained in the usual way, and the fire did not behave like that picture. The picture did not change as fast as the fire did. When the foreman, Wagner Dodge, saw that the fire was catching up with them, he ordered the men to drop their heavy tools; some did, some carried on with them. Then, with the fire less than a hundred meters behind, he lit a fire in the grass in front of them, stepped into the burnt area and called the men to follow. They did not.⁶

Weick calls such a moment a *cosmology episode*: people “suddenly and deeply feel that the universe is no longer a rational, orderly system”, and with the order goes the ability to act together.⁷ The order to drop the tools, which should have saved them, was for many of the men an order to stop being firefighters. The escape fire made no sense inside the picture they held. And the crew was a minimal organization, a group put together for the job; such groups are “susceptible to sudden losses of meaning”, because once the frame and the role structure go, no shared structure is left in which a new picture could be made quickly.⁸

Read with Hutchins — and this reading is mine, not Weick’s — the tools were more than weight. They were part of the means by which the men knew what they were and what they were doing, and dropping them removed a piece of the crew’s cognitive system at the moment it was most needed.⁹ The *Palau*’s team lost a central instrument and reorganized. The Mann Gulch crew lost its tools, its frame and its role structure at once, and disintegrated.

⁶Mann Gulch fire, 5 August 1949, Helena National Forest, Montana: 15 smokejumpers plus fire guard James O. Harrison; 13 died; foreman Wagner Dodge survived in his escape fire, Saltee and Rumsey through a crevice in the ridge; the order to drop tools. — Weick (1993), pp. 633, 635 (cosmology/minimal organizations; drop-tools); primary account Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.

⁷Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

⁸Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

⁹Author’s reading, not Hutchins’ own.

Courage and talent do not explain the difference. What each group had to hold its understanding in does: on the ship, overlapping knowledge, practiced procedures and a role structure that survived the failure of a part; in the gulch, a picture that broke and nothing left to build another in.

Both cases show the thing this book is about in its purest form. The failure at Mann Gulch lay in *understanding*, in the continuous, shared, fragile work of knowing what situation one is in. The deciding, the executing and the data all came after that work, and they could be no better than it. The two cases are an argument, not a comparison, and the reading of each is the book's own; what would count against it is an accident analyzed in comparable detail in which the actors' picture of the situation was accurate and shared, and the failure lay elsewhere — in equipment, resources or chance.

The navigator's problem

The Pacific navigators who crossed thousands of miles of open ocean without instruments give the same problem its positive form. They read birds, swells, wind and the color of the water, and understood their meaning. The navigator has no better map than the smokejumpers had; the navigator has no map at all in our sense. What the navigator has is a practice of continuous orientation: reading many weak cues, holding a provisional picture of where the canoe is, testing it against the next swell, revising it without drama. The knowing is done all the way across.

Hutchins, comparing Western navigation with navigation in the Caroline Islands of Micronesia, gives this a sharper form. The task is the same in both traditions, where are we and where is the destination, but the representations differ radically. For the Caroline Island navigator, "a canoe on course between islands is stationary and the islands move by the canoe"; the voyage is tracked as a sequence of *etak* segments, marked by a reference island

passing under successive star bearings.¹⁰ Knowing where one is depends on how one represents it, and choosing the representation is part of the navigator's craft.

Most management systems were built for the opposite assumption: that the situation can be established first, by analysis, research or a planning cycle, and that action follows from it. The assumption is reasonable where the environment changes slowly compared with the time it takes to understand it. It fails where the environment changes faster than it can be known. The book adopts that premise as its starting point, in the form it takes in the belief core of the practice the book grew out of:

Uncertainty is the normal condition of organisations, not a fault to be engineered away. The environment changes faster than it can be known, and acting in it changes the terrain — so knowledge of what works is produced in action rather than before it.¹¹

Understanding is never finished. If the environment changes faster than it can be known, every account an organization holds of its situation is already aging when it is written down. The strategy paper, the market analysis, the risk register: each is a snapshot of an understanding that was made at some point and has been decaying since. An organization's picture is always out of date to some degree. The open questions are how quickly it notices the gap and how cheaply it can revise.

Understanding is partly self-made. "Acting in it changes the terrain." When a company launches a product, customers react, competitors respond, regulators take note, employees draw conclusions. The situation the company then has to understand is in part the product of its own action. Weick calls this *enactment*: organizations do

¹⁰Hutchins (1995), ch. 2, p. 71, following Gladwin, T. (1970), *East Is a Big Bird*.

¹¹The Tautai company belief core ("probes with stop conditions rather than plans with milestones"; "much of the work worth doing is not in a race"; delegation only as deep as the capacity to evaluate what comes back).

not simply meet an environment that is out there; they help produce the environment they then try to read.¹² The fire at Mann Gulch was real. Which fire the crew believed it was fighting was a matter of the picture they had built, and that picture shaped what they did and where they were when the fire turned on them. And the environment contains more than terrain: customers, competitors and regulators are organizations too, building pictures of their own, in part of us.

Understanding is distributed. Nobody on the *Palau* held the whole picture, and nobody at Mann Gulch did either. Organizations are like this all the time, only more so. The customer's changing language is heard by a service agent; the recurring technical workaround is known to an engineer; the budget pressure is felt in finance; the shift in regulation is read by legal. What the organization "knows" is spread across people, documents, systems and routines, and none of them holds it whole.¹³

Put the three together and the problem is clear. An organization has to act on an understanding that is always provisional, partly produced by its own action, and held by no one in particular. How well it adapts depends less on how fast it decides than on how well it keeps that understanding alive: how early it notices that the picture is wrong, how well it can build a new one together, and how much of what it learns from each round stays available for the next.

How the book is built

The book follows understanding from perceiving to learning, then drops the simplification that one actor does it, then turns to what it produces. Part I (chapters 1–2) states the problem of knowing and takes up perceiving: attention, and the signals at the edges. Part II (chapters 3–6) is about understanding the situation: being aware,

¹²Weick, K. E. (1995). *Sensemaking in Organizations*. Sage. Enactment as a property of sensemaking; equivocality vs uncertainty. (pp. 27, 30–31 (enactment; equivocality vs uncertainty)).

¹³Hutchins (1995), introduction (chapter summaries).

making sense when the picture breaks, keeping the model current, and choosing, where the question of speed is settled. Part III (chapters 7–10) asks where understanding is held, when groups understand more and less than their members, how structure carries what a unit faces, and whose account counts. Part IV (chapters 11–14) turns to what understanding produces: learning loops and how knowledge is created and lost, the path from knowledge to judgment, making sense with machines, and the design of an understanding organization. Appendix A maps the book's terms onto *The Tautai Principle*, the operating model the book grew out of, and says where the book departs from it; Appendix B is a glossary. Readers who want the theory first can start with Parts II and III and return. Readers who come with a design task can start at chapter 14 and follow its references back. The argument of every chapter is in the outline.

Why more data does not solve it

The obvious objection is that we have never had more data. Organizations measure more than ever, dashboards update in real time, and machine learning finds patterns no analyst would see. Surely the problem of knowing one's situation is shrinking.

Herbert Simon answered this objection in 1971, before most of today's data existed: "a wealth of information creates a poverty of attention".¹⁴ Information consumes the attention of its recipients. The more of it there is, the scarcer the capacity to notice what matters in it. More data raises the demand for the work this book is about.

There is an organizational version of the same point, which practitioners know as the measurement trap: the numbers are accurate, and by the time a problem shows up in quarterly revenue or a retention score, the change that

¹⁴Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press. The poverty-of-attention sentence and the design rule are in the OCR'd scan (printed pp. ~40–42); see chapter 1.

produced it is months old.¹⁵ The early signals of change are rarely in the metrics: a new objection in sales conversations, a pattern in support tickets, a workaround that spreads. The metrics were designed to track a situation that was understood at the time they were designed. A measurement system encodes an old picture. It tells you how far reality has moved along the axes you expected, and little about whether the axes have changed.

A second reason is less obvious. Data does not say what it means. Two managers can look at the same drop in a satisfaction score and see a temporary effect of a price change, the first sign of a competitor's success, or noise. They disagree about the picture the number is read into, and the number cannot settle that. Where the situation is familiar, the right picture is usually obvious and analysis can do the rest. Where the situation is new, several pictures are plausible at once, and the work is to find one that is good enough to act on and cheap enough to revise. That work is sensemaking, in Weick's sense. Data is its raw material.

Understanding, taken apart

"Understanding the situation" covers several different activities. In everyday language they blur into each other; in organizations they are often done by different people in different places at different speeds, and they fail in different ways. The book keeps five apart, each with an individual and a collective form.

Perceiving is attending and noticing. Attention decides what gets looked at at all. It is limited in individuals and in organizations, and agendas, dashboards, meeting formats and escalation paths decide, mostly by accident, what an organization is able to notice. Knowing what to ignore is as much a design question as knowing what to watch. Noticing picks up early, ambiguous cues of change, weak signals. Such cues almost always exist; what varies

¹⁵Mathis, K. (2025). *The Tautai Principle*, ch. 4, p. 72. Paraphrased in chapter 1.

is whether they travel from the people who see them to the people who can act on them. Chapter 2.

Understanding is being aware and making sense. Being aware means holding an accurate enough picture of the current situation and where it is heading. Mica Endsley's model of situation awareness, published in 1995, distinguishes three levels: perception of the relevant elements, comprehension of what they mean together, and projection of their state in the near future.¹⁶ Situation awareness assumes that there is a situation to be perceived and that the categories for perceiving it are available. Making sense is what people do when that assumption fails: when cues do not fit, when several interpretations are plausible, when something happens that should not have happened. Weick's sensemaking constructs a plausible, shared account that is good enough to act on, and uses action to test it. Chapters 3 and 4; the model that is kept current through both, chapter 5.

Choosing is selecting and committing. For experienced people, understanding already contains the choice. Gary Klein found that fireground commanders made "around 80 percent of their decisions in less than one minute": they "did not have to compare options" and "could come up with a good course of action from the start", because recognizing the situation as a type delivered a workable action with it.¹⁷ Most choosing in organizations is like this: continuous, tacit, part of understanding. The book calls it *action selection*. A *decision* is something else, an explicit, accountable commitment of resources that others can rely on and later check. Organizations need both, and they fail differently. Action selection fails when the model behind the recognition is wrong; decision fails when commitments are made without a record, so that nothing can later be learned from them. Chapter 6.

¹⁶Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. Chapter 3 gives the model in full.

¹⁷Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press, pp. 12, 17 and 24 (printed page = PDF page – 24).

Acting turns a commitment into an intervention, and every action is also a probe. It changes the terrain and produces new cues, and what it reveals about the picture is as important as what it achieves.

Learning is what the organization takes from the consequences of its action. Some of it changes the organization's picture of the world; some changes its governing assumptions; some settles in people as knowledge, skill and, over time, judgment. Much of what matters here is tacit. Ikujiro Nonaka, writing in 1994 and following Michael Polanyi's "We can know more than we can tell", distinguishes explicit knowledge, "transmittable in formal, systematic language", from tacit knowledge, which is "hard to formalize and communicate" and "deeply rooted in action, commitment, and involvement in a specific context".¹⁸ Organizations are good at managing what is already explicit, the reports, dashboards and knowledge bases, and poor at the rest: sharing what people know from experience, putting half-known things into words, and turning written knowledge back into practice. Nonaka and Hirotaka Takeuchi's model of knowledge creation is the core theory of Part IV. Judgment in turn is what lets people notice the next cues, make sense of them, and recognize the action. Chapters 11 and 12.

Each activity fails in its own way and needs its own remedy, which is why the distinctions matter in practice. An organization that does not attend to a market will not notice signals from it, however good its analysts are. An organization that notices signals but has no forum in which they can be interpreted will collect anecdotes. An organization that is aware of its current situation but cannot make sense of a surprise will do what the Mann Gulch crew did: keep executing a plan that no longer fits. An organization whose experts recognize the right action but whose decision process demands option comparison will be slow for no gain; one that commits without records will

¹⁸Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4. (public PDF, pp. 14–21). Nonaka, I., & Takeuchi, H. (1995), *The Knowledge-Creating Company*, Oxford University Press.

not learn. An organization that makes sense well but does not learn will make the same sense again next year, at the same cost.


Activity	Question	Typical failure	Chapter
Perceiving: attending, noticing	What do we look at, and what is starting to change at the edges?	Data overload; attention captured by the familiar; signals filtered out before they reach anyone who can act	2
Understanding: being aware	What is happening, what does it mean, what comes next?	Tunnel vision; errant mental models	3
Understanding: making sense	What is going on here, and what story lets us act?	Premature closure; parallel realities across functions	4, 5
Choosing: selecting, committing	What action does this situation call for, and what do we commit to?	Recognition on a wrong model; decisions without records; option comparison where recognition would do	6

Activity	Question	Typical failure	Chapter
Acting	What does the action achieve, and what does it reveal?	Actions that are not read as probes	6
Learning	What did our action teach us, and what do we keep?	Same lesson learned repeatedly; nothing unlearned; solutions not saved in the system	11, 12

The title of this chapter needs one qualification, then. “Understanding comes first” names a dependency: choosing, acting and learning can be no better than the understanding they rest on. It says nothing about a sequence of stages. For experienced people in familiar situations, understanding and the action arrive together, and the separation into stages is something novices and novel situations need.

One loop, not a pipeline

Laid out on a page, the five activities fall into a line, and a line is how most process diagrams would draw them. Organizations do not run in a line, and the book will not treat them as if they did. The activities run in parallel, at different speeds, and they feed back into each other. What one learns changes what one attends to. Acting produces new cues. Making sense of a surprise changes what counts as a signal. The more useful picture is a loop, and a loop that is held by more than one person:



The loop has an inner and an outer circle. The inner circle, from cue to account to action and back to new cues, is the one John Boyd described as observe–orient–decide–act, and it can turn fast. The outer circle, from action to learning to changed capability to changed attention, turns slowly, over months and years. Organizations often invest in the inner circle, in faster reporting and shorter decision paths, and neglect the outer one. They get quicker at acting on a picture they no longer improve.

The loop also settles where speed belongs. Much of the management literature on loops of this kind presents them as a race: the organization that cycles faster sets the tempo. That holds where a race is the frame, and much of the work worth doing is in no race. What the loop must deliver is valid orientation; a fast loop on a wrong picture only takes the organization further in the wrong direction, faster. Klein’s commanders were fast because their models were good; nobody had pressed them for speed. Chapter 6 argues that speed follows from model quality and cannot be pursued on its own.

And the loop is drawn inside a box for a reason. None of the five activities is done by an organization the way a person does them. Cues are seen by one person and reported to another; an account is built in a meeting and written into a document; a decision is recorded, or not; what was learned is kept in a procedure, in a story, or in nobody. The box — people, artifacts, forums, structure —

is not the container of the loop. It is what does the loop, and Part III is about it.

The model behind the loop

One element of the loop needs a name before the book goes further. The box at the bottom of the diagram, knowledge, skill and judgment, is where the organization keeps its working account of the world and of itself. Kenneth Craik suggested in 1943 that thought works with a “small-scale model” of external reality, and much of what followed in cognitive science rests on that idea.¹⁹ This book calls it the *internal model* and uses the term for the account an organization acts on, whether it sits in heads, in documents and dashboards, or in the shape of the structure itself.

The internal model runs through all five activities. It decides what is attended to. It sets which distinctions can be made at all, and therefore what can count as a signal: Ross Ashby, the cybernetician who formulated the law of requisite variety, noted that “a set’s variety is not an intrinsic property of the set: the observer and his powers of discrimination may have to be specified”.²⁰ It carries comprehension and projection, which is why Endsley and Debra Jones call mental models “the key enablers” of the higher levels of situation awareness.²¹ It is the frame that sensemaking repairs when it fails. It is what recognition runs on when an expert sees the action in the situation. It is run forward in planning and backward when people explain what has happened. And it is what learning changes.

Stafford Beer, whose Viable System Model the book uses in Part III, put the organizational version plainly in

¹⁹Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.

²⁰Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

²¹Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press, pp. 21–22.

Brain of the Firm. The part of an organization that looks ahead “must contain a model — some model — of the corporation”. That model “may be disseminated, in separate chunks of cognition around the firm, rather than being cohesive and well formulated. Certainly, no one may think of it, or refer to it, as a model, but it must be there.” On the same page he draws the consequence that matters most for this book: “a corporate plan must continuously abort, for unless that plan is continuously aborting, how shall sanity (the match between the world and the firm’s notion of it) be preserved?”²² Having a model and following a plan are different things. The navigators of the Caroline Islands had a model, the *etak*.

The book uses the term in two senses and keeps them apart. In the *loose* sense, every organization acts on some account of its situation, and the only questions are where it lives, who drew its distinctions and how it is revised. That sense is uncontroversial. In the *strong* sense, a regulator needs a model of what it regulates, and steering is no better than the model. The cybernetic tradition has often claimed that this is proved, citing Conant and Ashby’s 1970 paper “Every good regulator of a system must be a model of that system”. The paper does not prove its title, as its authors concede in the text; the rigorous form of the claim is Francis and Wonham’s internal model principle of control theory, which holds for control problems with stated dynamics, and an organization is not obviously one.²³ Where the book uses the strong sense it argues for it and says so. It claims no proof.

A second caution is Weick’s. He tells of a reconnaissance unit lost in the Alps that found its way back with a map of the Pyrenees: “when you are lost, any old map will do”. What saved the soldiers was that they kept moving,

²²Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

²³*Punch and Judy*, ch. 3 “The Law We Quote Is Not the Law Ashby Wrote”, section “The theorem that is not a theorem”, on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors’ concessions at p. 96. Beer’s “cybernetically proven fact” at *The Heart of Enterprise*, p. 114.

“kept noticing cues, and ... kept updating their sense of where they were”.²⁴ A model does not have to be accurate to start with. It has to be good enough to act on and open to correction by what the action produces. Chapter 5 develops the internal model as the core of orientation and settles where the book stands between Weick’s map and Beer’s regulator; chapter 9 develops it as the core of the Viable System Model’s System Four.

Who is doing the understanding?

Each time this chapter has said that “the organization” notices, understands or learns, it has used a figure of speech, and the figure has still to be earned.

At one level, all understanding happens in individual minds. People attend, perceive, interpret, remember. Individual cognition has well-studied limits: working memory, selective attention, the fast and slow modes Kahneman described in 2011. A good part of what goes wrong in organizations can be traced to them.²⁵ Parts I and II take this level seriously.

But individual minds are not where most organizational understanding is held. Hutchins’ claim about the *Palau* is more than a metaphor because it comes with a mechanism. Thinking, in his terms, is “the propagation of representational state across a variety of media”.²⁶ A bearing seen through the sight of an instrument reaches the bearing recorder as “strings of spoken digits”, becomes an entry in the bearing log, and becomes a line on the chart.²⁷ Each step can be observed, and each can fail. That turns “the organization thinks” into a set of questions one can actually ask: where is this computation done, which step is carried by a person and which by an artifact, and where does the picture get lost on its way? Two consequences

²⁴Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 54–55.

²⁵Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux. The scan carries no printed folios, so no page is asserted; the two-modes account is his.

²⁶Hutchins (1995), introduction (chapter summaries).

²⁷Hutchins (1995), ch. 4, p. 226.

follow, and Part III develops both. Such a system can be smarter or dumber than its members, depending on how its communication is organized; “it is not the case that two or more heads are always better than one”.²⁸ And its artifacts and structures are part of the thinking, so changing a form, a dashboard, a handoff or a role changes how the organization thinks.²⁹

Organizations are cognitive systems in exactly this sense. Their memory sits partly in people and partly in documents, tickets, budgets and the knowledge of who knows whom; their attention is routed by calendars and dashboards; their inference is shaped by what is easy to represent and what is safe to say. An operating model is, among other things, a thinking model: it decides where the organization’s working picture of a situation lives, who can challenge it, and how it crosses from one unit to another.

And organizations are recursive. A team faces its own environment; so does a business unit; so does the company as a whole. In the language of Beer’s Viable System Model, each viable unit has its own outward-looking function and its own balance between running what exists and adapting to what is coming.³⁰ Understanding therefore has to be produced at every level that faces its own environment, and one picture for the whole company hides exactly what its parts need to see. Beyond the company there are networks, markets and ecosystems that face environments of their own.

Three levels, then, stay in view throughout: the person, the group, the organization and its units. The question is the same at each. Where is the picture of the situation held, who may change it, and how?

²⁸Hutchins (1995), introduction; ch. 5, p. 253; ch. 4, pp. 221 and 227.

²⁹Hutchins (1995), ch. 3, pp. 154 and 169.

³⁰Beer, S. (1979). *The Heart of Enterprise*; Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley. (Beer 1979, *The Heart of Enterprise* pp. 5 and 118 — Recursive System Theorem).

The theories the book uses

No single theory covers the loop. The book draws on several, each for one question, and says which one it is using at any point.

Question	Theory	Chapter
How is understanding built, and why does it collapse?	Karl Weick — sensemaking	4
What does “knowing what is going on” consist of?	Mica Endsley — situation awareness	3
How is the picture kept current, and what is a model for?	John Boyd — orientation; Craik, Beer, Klein — the internal model	5
How does understanding become action?	Gary Klein — recognition-primed decisions; Boyd — decision as hypothesis	6
Where is understanding held, and how does it travel?	Edwin Hutchins — distributed cognition	7
When does a group understand more than its members?	Anita Woolley, Amy Edmondson, transactive memory	8
Which structures could carry what a unit faces?	Stafford Beer — Viable System Model; Ashby — requisite variety	9

Question	Theory	Chapter
Whose account counts?	Dennis Gioia, Sally Maitlis — sensegiving	10
How is new knowledge created and spread?	Ikujiro Nonaka and Hirotaka Takeuchi — SECI	11–14
How does learning become judgment?	Hubert Dreyfus, Anders Ericsson, Hutchins, Michael Polanyi	12
Which situations need which kind of understanding?	Dave Snowden — Cynefin	1, 6

These theories do not agree with each other, and the book does not pretend they do. Weick is suspicious of models and trusts action; Beer and Hutchins are built on representations. Nonaka starts from the individual — knowledge “is created by individuals” and the organization amplifies it into something shared; for Hutchins the system of people and artifacts is what thinks.³¹ Klein trusts an expert’s recognition; in complex situations, as this book reads Cynefin, that recognition can be confidently wrong. The book treats these disagreements as productive and settles none of them early. Where it takes a side, it says so, and it marks its own readings as its own.

When understanding needs sensemaking

Not every situation needs sensemaking. Where cause and effect are known and stable, analysis and good practice are enough; bringing a team together to ask “what is going on here?” about a routine delivery problem is waste.

³¹Nonaka (1994), p. 17. Hutchins (1995), introduction.

Where cause and effect are knowable but not obvious, expert analysis is the right tool. Dave Snowden's Cynefin framework, as Snowden and Boone set it out in 2007, draws the line between the complicated and the complex: in the complex domain, patterns emerge from interaction and can be understood in retrospect but not predicted, and the appropriate response is to probe, sense and respond rather than to analyze and plan.³² Analysis assumes the situation can be broken into parts, studied and understood through data; sensemaking applies where observing changes what is observed and several interpretations coexist.

Three kinds of occasion call for sensemaking in organizations:

- **Surprise.** Something happens that should not have happened, or something expected fails to happen. The existing picture has been contradicted.
- **Equivocality.** The cues are there, but they support several plausible readings at once. More information does not help, because nothing is missing; there are too many meanings. Weick distinguishes this from uncertainty, where information is missing and can in principle be gathered.³³
- **Novelty.** The situation belongs to no category the organization has. New technologies, new regulation, a merger, a crisis.

Organizations are poor at recognizing these occasions. Surprise is explained away as noise; equivocality is resolved prematurely by whoever speaks first or has the most authority; novelty is forced into the nearest familiar category. Mann Gulch had all three at once, and the crew's frame absorbed each of them until it could not. Chapter 6

³²Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76. Klein's and Snowden's positions are as summarized in the theories table of chapter 1; that they "disagree about how often each case occurs" is my reading, not a documented exchange.

³³Weick, K. E. (1995). *Sensemaking in Organizations*. Sage. Enactment as a property of sensemaking; equivocality vs uncertainty. (pp. 27, 30–31 (enactment; equivocality vs uncertainty)).

returns to Cynefin, because the domain of a situation also decides what recognition can be trusted to do.

What this book claims — and what would count against it

A model that cannot lose is a marketing device. The book therefore states, for each claim, what would count against it and where such evidence exists or would have to be found. The overall claims are these:

1. **Understanding, not decision speed, limits adaptation in complex environments.** *Against it:* organizations that invest in shared understanding, in forums, representations and time for orientation, adapting no better than organizations that invest only in faster decision paths or more data, in comparable complex settings. *Where to look:* the high-reliability literature, the accident analyses, the studies of team learning.
2. **Understanding is produced in a loop that includes action and learning, and has no phase of its own before action.** *Against it:* organizations that separate understanding from action, plan first and then execute, adapting as well in complex settings as organizations that use action to test their picture. *Where to look:* naturalistic decision research; studies of probing and experimentation in strategy.
3. **Organizational understanding is distributed across people and artifacts, and can be designed.** *Against it:* changes to representations, artifacts, forums, handoffs and structures leaving what an organization notices and how it interprets unchanged. *Where to look:* transactive-memory and distributed-cognition studies; the effects of forum and reporting design on what surfaces.
4. **Learning closes the loop only when it changes capability; records on their own do not close it.** *Against it:* organizations that keep lessons-learned records without changing practice or skill learning as well as those that change practice and skill. *Where*

to look: knowledge-management evaluations; the experience-pipeline studies of Part IV.

The book cannot prove these claims. It argues for them from the literature and from practice, and it states them so that a reader, or a client, can check them against their own experience and say where they fail.

The fourth claim is where the *Palau* comes back. In twenty-five minutes its navigation team did what the first three claims ask for: it kept a picture of where the ship was alive while acting on it, fix by fix, and it held that picture in a system of people, charts and instruments rather than in one head. Then the ship anchored, and “the solution ... was not saved in the system”.³⁴ The team had done the understanding and kept none of it. What it takes to keep it, from the first cue at the edge of attention to the judgment that recognizes the next one, is what the rest of the book is about.

References

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
- Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.
- Weick, K. E. (1985). Cosmos vs. chaos: Sense and nonsense in electronic contexts. *Organizational Dynamics*, 14, 50–64.
- Gladwin, T. (1970). *East Is a Big Bird*.
- The Tautai company belief core, section “What this is for”.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press.

³⁴Hutchins (1995), ch. 8, p. 351.

- Mathis, K. (2025). *The Tautai Principle*.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Polanyi, M. (1966). *The Tacit Dimension*.
- Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.
- Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- *Punch and Judy*.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
- Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica*, 12(5), 457–465.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley.
- Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76.

Chapter 2

Chapter 2 — Perceiving: attention and the signals at the edges

A word in italics

On page 352 of *The Heart of Enterprise*, published in 1979, Stafford Beer sets one word in italics and stops to say why. The word is *model*. “*Model* is italicized, because it must not be confused with the variety in the information that happens to be flowing”, he writes, and the example he gives is a family: parents who cannot recognize the symptoms of their son’s drug use.¹

Beer does not say that the parents had too little information. The reading I take from the page, and it is mine, is that the information was flowing past them all the time. What they lacked was a model in which what they saw could mean what it meant. A symptom is only a symptom for someone who holds the distinction it belongs to. Without that distinction it is a mood, a late night, a phase.

That page is the whole of this chapter in small. Chapter 1 named perceiving as the first of five activities and said that attention is limited in people and in organizations. This chapter asks what limits it. What does an organization look at, who decided that, and how does an early, ambiguous cue travel from the person who sees it to the person who could act on it? The answer runs through three filters. The first is the one Beer’s parents show: the model decides what can be seen at all. The second is the

¹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

set of channels a cue has to travel. The third is silence, the people who see and do not say. None of the three is a matter of capacity in the first place, and each of them was designed by someone, usually without meaning to.

Attention is designed attenuation

Start with the person, because the organizational limits are built from the individual ones and then made worse.

Herbert Simon's sentence from 1971, quoted in chapter 1, is the compact form: information consumes the attention of those who receive it, so a wealth of the one produces a poverty of the other.² Attention is the scarce input of perceiving. Daniel Kahneman's account of two modes of thinking gives the scarcity a shape: a fast, automatic mode that runs on pattern and habit, and a slow, effortful mode that can check, doubt and compare, and that tires.³ The slow mode is the one that notices when something fails to fit, and it is also the one in short supply; under load the fast mode takes over without announcing that it has. Working memory, the space in which a person can hold and compare a few things at once, is small.⁴ Daniel Simons and Christopher Chabris showed in 1999 how far this goes: people who are busy counting one thing can fail to see something large and strange passing in front of them, and several did not believe the event had happened until the

²Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press. The poverty-of-attention sentence and the design rule are in the OCR'd scan (printed pp. ~40–42); see chapter 1.

³Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux. The scan carries no printed folios, so no page is asserted; the two-modes account is his.

⁴Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux. The scan carries no printed folios, so no page is asserted; the two-modes account is his.

videotape was replayed for them.⁵ Their eyes were open; their attention was spent on the count.

Now the organization. Its attention is allocated by things that nobody usually thinks of as attention devices. The monthly business review looks at what its template asks for. The dashboard displays the measures that somebody chose when it was built. The ticket system routes what its category list has a name for. The escalation path carries what someone dared to escalate, to the level that the rules of the path allow. The management calendar decides which questions get an hour and which get none. Each of these was designed, once, by a person with a purpose. The sum of them was designed by nobody. It is the organization's attention, and it decides, before any person has looked at anything, what can be looked at.

Beer's word for what these devices do is attenuation: they reduce the variety of what comes in to what the receiver can handle.⁶ The word is useful because it removes a false alternative. An organization does not choose between filtering and not filtering. It cannot look at everything, and every agenda, report and role is already a filter. The choice is between filters that were set on purpose, and are known, and filters that accumulated and are not.

The collective form of attention is the forum. A meeting is an attention device with people in it, and the list of who is in the room decides which models are in the room. A steering committee made up of the heads of functions attends to what the heads of functions can distinguish; the categories of the frontline are represented, if at all, by the one person who was invited to present and left after the item. That is an observation about where the partitions sit, and no complaint about hierarchy. The organization's attention is the sum of its forums, its templates, its systems

⁵Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059–1074. p. 1068 (54% noticed the unexpected event, 46% did not — the figures are pooled across all conditions, not the gorilla condition alone), p. 1072 (“several did not believe the event had happened until the videotape was replayed”).

⁶Beer (1979), on attenuators and amplifiers of variety. The terms are Beer's.

and its paths, and a designer who wants to know what the organization can notice can read it off those, before asking anyone what they think they watch.

Requisite variety as an argument, not a figure

The cybernetic tradition has one law about this, and it is easy to misuse, so the chapter sets it out before the book relies on it.

Ross Ashby wrote, in chapter eleven of *An Introduction to Cybernetics*, that “only variety in R can force down the variety due to D; variety can destroy variety.”⁷ R is the regulator, D the disturbances it faces. The sentence most people quote is a different one: “only variety can absorb variety”. That is Beer’s compression, from *Brain of the Firm*, and it is the better sentence, which is why it traveled and Ashby’s did not. What it left behind on the way were the conditions under which Ashby’s result is exact.⁸ *Punch and Judy*, the book on the Viable System Model that this book follows on all cybernetic questions, sets out three of them.

First, the law is information theory applied to regulation. Ashby ties it to channel capacity: “R’s capacity as a regulator cannot exceed R’s capacity as a channel of communication”.⁹ A regulator that cannot receive a distinction cannot act on it. Second, the law is a priori. In Ashby’s words “it has nothing to do with the properties of matter”; it holds for any regulator of any kind, which is exactly why it licenses an argument about whether an arrangement *could* work and never a measurement of whether it does.¹⁰ Third, the one careful attempt to turn it into a

⁷Ashby (1956), p. 207 (section 11/7), pagination as in *Punch and Judy*, ch. 3. “Only variety can absorb variety” is Beer’s compression (*Brain of the Firm*, 2nd ed., p. 308).

⁸*Punch and Judy*, ch. 3, “Two sentences and one verb” and note 2, on Beer, *Brain of the Firm* (2nd ed., 1981), p. 308.

⁹*Punch and Judy*, ch. 3, note 4, on Ashby (1956), pp. 209–211 (section 11/11).

¹⁰*Punch and Judy*, ch. 3, note 5, on Ashby (1956), p. 208 (section 11/10).

count, Pffiffner's in 2017, ended with a judgment of adequacy and no number.¹¹ This book therefore reports no variety figures anywhere, and it quotes Ashby's sentence, not Beer's.

What the law does give this chapter is a way of asking a design question. Take a business unit facing a market that changes week by week, and the one channel through which that unit's picture of the market reaches the level that decides its budget: a quarterly review with a fixed template. One does not need to count states to see that the channel cannot carry what the market presents. The argument is about what could in principle pass, and it is the argument organization designers already make when they say that a structure "cannot see" something. Ashby lets them say it with a reason. The law says nothing about whether what does pass is worth seeing. That is the next question.

The model bounds what can be noticed

Some eighty pages before the sentence about variety destroying variety, in section 7/5, Ashby had said something that the tradition quotes far less. "A set's variety is not an intrinsic property of the set: the observer and his powers of discrimination may have to be specified if the variety is to be well defined."¹² Variety is a count of states that are treated as different, and somebody has to have decided which states those are.

Punch and Judy calls this the partition problem and treats it as the point where an observer enters the model, and with the observer a politics: whoever draws the partition decides what the regulator can distinguish, and everything that falls between two of its lines is, for that regulator,

¹¹*Punch and Judy*, ch. 3, section "What the law becomes when someone has to measure it", on Pffiffner (2017), p. 165; the a-priori character of the law, Ashby (1956), p. 208 (section 11/10). (P&J)

¹²Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

one thing.¹³ In an organization the partitions are easy to find once one looks for them. The segments in the market model. The categories in the support system. The line items in the budget. The taxonomy of risks in the register. The list of reasons a customer may give for leaving, as the cancellation form offers them. Each list was drawn by someone, at some time, for some purpose, and each is now a statement about what the organization can perceive. A customer who leaves for a reason the form does not offer picks the nearest one, and the organization learns that the nearest one is rising.

This is the mechanism behind Beer's parents. Beer draws the consequence in the same book, in two places. Early on, in dialogue: "we can only control a system to the extent that we understand it. ... all we're controlling is the projection of the model-in-our-heads."¹⁴ And later, as a claim: "the regulator is as good as its contained model of the system regulated, and no more".¹⁵ Beer calls this a "cybernetically proven fact". Chapter 1 said why the book cannot repeat that: the proof he leans on, Conant and Ashby's, does not prove its title, and the book argues for the strong sense of the model where it uses it rather than claiming a theorem.¹⁶ But the weaker claim needs no theorem. Whatever the organization's model cannot distinguish, the organization's attention cannot be directed at. The parents were attending, with a model that had no place for what they saw.

The same point answers a question chapter 1 left open: why more data does not solve the problem of knowing. Data arrives already partitioned. It comes through the

¹³*Punch and Judy*, ch. 3, section "Who counts the states"; boundary in ch. 5, transducer in ch. 25.

¹⁴Beer (1979), p. 80.

¹⁵Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

¹⁶*Punch and Judy*, ch. 3 "The Law We Quote Is Not the Law Ashby Wrote", section "The theorem that is not a theorem", on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors' concessions at p. 96. Beer's "cybernetically proven fact" at *The Heart of Enterprise*, p. 114.

categories of whatever system collected it, and those categories are the old model, written into the collection. More data sharpens the picture along the axes the model already has. It cannot add an axis, because nobody asked for that axis when the fields were defined.

Weak signals are filtered, not rare

Igor Ansoff gave the term *weak signals* its currency in 1975, in an argument that discontinuities announce themselves early and vaguely and that a strategy could respond to them before they hardened into facts.¹⁷ Much horizon-scanning practice has grown from that idea, and much of it rests on an assumption the term itself invites: that weak signals are faint and few, so that the task is to find them. The account in this chapter says something different. The signals are usually there. What makes them weak is the organization.

The first filter is the model, and it has just been described. A cue that fits no distinction in the model is not weak. It is invisible: it arrives as noise, or as a slightly odd instance of something familiar. This is the sharpest explanation the book has for the pattern practitioners know well, in which the change that later seemed obvious was in the data all along, unlabeled. The measurement trap of chapter 1 is a special case: the metrics encode an old picture, and a change that the picture has no axis for shows up, if at all, as unexplained movement on an axis that was built for something else.¹⁸

The first filter has an individual form as well. Hutchins, in his study of the navigation team, took confirmation bias as his example of a property that people carry into a group: “a propensity to affirm prior interpretations and to discount, ignore, or reinterpret evidence that runs counter to an already-formed interpretation”.¹⁹ A prior interpretation

¹⁷Ansoff, H. I. (1975). Managing strategic surprise by response to weak signals. *California Management Review*, 18(2), 21–33.

¹⁸Mathis, K. (2025). *The Tautai Principle*, ch. 4, p. 72. Paraphrased in chapter 1.

¹⁹Hutchins (1995), ch. 5, p. 239.

is a model. The person who holds it does not refuse the contrary cue; the cue is read into the interpretation and comes out as support for it, or as an exception. Diane Vaughan's study of the *Challenger* launch decision describes the organizational version, which she called the normalization of deviance: an anomaly is noticed, assessed, and reclassified as acceptable, and the next anomaly is assessed against the new baseline, so that the distance from the original standard grows one reasonable step at a time.²⁰ Each step moved a line in the partition. The signal did not weaken. The model moved to absorb it.

The title of this chapter speaks of the edges, and the word means two things. There are the edges of the organization: the counter, the call, the site, the boundary with a supplier or a regulator, where the environment touches people who are not in the forums. And there are the edges of the model: the margins of its categories, where a case sits that is almost one thing and almost another, and is filed as one of them. Weak signals appear at both. A cue that arrives at the edge of the organization has the longest way to travel; a cue that sits at the edge of the model has the smallest chance of being seen as anything at all. The two edges often coincide, because the people at the boundary are the ones whose categories differ most from those of the model in use at the center. That is the reason to look there, and it is also the reason it is hard: what those people see is described in a vocabulary the forums do not use.

The second filter is channels, and the third is silence. Both are the subject of the next section, because both are social, and that is where the practical work of this chapter lies. But see now what follows from putting the three together. If signals are filtered rather than rare, then missing them calls for an examination of the filters: which distinctions the model makes, which channels a cue has to cross, and who is in a position to see something and has reason to keep quiet. Karl Weick and Kathleen Sutcliffe's studies of organizations that operate under high hazard with few

²⁰Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.

accidents describe two habits that read as counter-filters: a preoccupation with failure, which treats small failures as signals rather than as closed cases, and a reluctance to simplify, which resists reducing the picture to the few categories that make it manageable.²¹ Both work on the first filter: they keep the partition open.

Channels, amplification and silence

A cue that has been seen still has to travel. In most organizations it travels through people, and each person it passes re-codes it: the service agent's impression becomes a note, the note becomes a line in a weekly report, the line becomes a bullet on a slide, the slide is read by someone who was not in the conversation. Chapter 7 follows this path in detail, with Hutchins' account of how a bearing moves from an instrument to a chart. Here it is enough to see that every step is a filter of the second kind, and that the filter can be set well or badly. A channel that requires the cue to be classified before it can be reported has already applied the model. A channel that requires it to be significant before it can be escalated has asked the person who saw it to do the interpreting alone.

A channel built for early cues looks different from one built for reports. It accepts an impression that has no category yet, in the words of the person who had it. It does not ask that person to say what the impression means or how important it is, because those are the questions the organization has not yet been able to answer, and asking them at the point of entry is how the first filter gets applied a second time. It has a receiver who is expected to read such impressions, in the slow mode, at a rhythm that fits how fast the environment moves. And it has a return path, so that the person who raised the cue learns what became of it, since a channel on which nothing is ever heard back is a channel that stops being used. Time enters here in a way that chapter 14 takes up as a design variable: a cue is

²¹Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the Unexpected: Sustained Performance in a Complex World* (3rd ed.). Wiley.

early only relative to the rhythm at which the organization looks. A monthly forum makes every cue at least a month old before anyone with authority sees it, and the filters that decide what reaches that forum are themselves revisited only when someone remembers to.

Beer's word for the opposite of attenuation is amplification, and the two belong together: a cue from the edge of the organization is faint at the center unless something on the way makes it louder.²² Amplification is not the same as more traffic. Hutchins' simulations of small groups showed "that more communication is not always in principle better than less", and that richer communication can produce worse properties at the level of the group.²³ What a weak signal needs is a channel on which it is not competing with everything else, and a receiver who can give it the slow mode of attention it requires.

Then there is silence. Elizabeth Morrison and Frances Milliken gave it a name in 2000, organizational silence, for the condition in which employees who know about problems withhold what they know, because saying it seems futile or dangerous.²⁴ Silence is the third filter and in practice often the decisive one. The people who see a weak signal first are, by the nature of the thing, the people closest to where it appears: at the counter, on the line, on the call, in the code. Beer knew where such people sit. Asking who really decides to buy, he answered: "It is the little group of employees, possibly very junior, which understands the need".²⁵ The same little group is the one that sees the need change. Whether what they see reaches anyone depends on whether raising it is safe, whether it is expected, and whether anything ever happened the last time someone raised something.

²²Beer (1979), on attenuators and amplifiers of variety. The terms are Beer's.

²³Hutchins (1995), ch. 5, pp. 239 and 252.

²⁴Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725. p. 706 (organizational silence defined), p. 708 (a climate of silence: futile and/or dangerous).

²⁵Beer (1981), p. 232 (McCulloch's redundancy of potential command).

This is why the chapter treats detecting weak signals as a social activity rather than as the job of a scanning department. A scanning unit can read the environment that is already partitioned into the categories the unit was given. It cannot be at the counter. *Punch and Judy* makes the structural version of this argument: Beer's System Four, the part of an organization that looks ahead, is a modeling function rather than "a pair of eyes", and teaching it as a scanner removes the modeling from what is taught.²⁶ The scanning that matters is done by many people, in the course of their work, with whatever model they hold. And the model they hold is largely tacit. Ikujiro Nonaka, following Michael Polanyi, distinguished knowledge that is "transmittable in formal, systematic language" from tacit knowledge, which is "hard to formalize and communicate" and "deeply rooted in action, commitment, and involvement in a specific context".²⁷ The experienced agent who hears a new kind of objection on a call, and knows it is new, is using tacit knowledge to recognize a weak signal. The two are not the same thing. The signal is in the world. The knowledge that lets it be seen as a signal is in a person who has taken a great many calls. An organization that loses such people, or does not listen to them, loses the capacity to receive the signal, whether or not the signal is still there.

Attention and action selection

There is a symmetry here that the book will use again. Attention selects among what comes in; something has to select among what could go out. *Punch and Judy* calls action selection "the function the diagram omits" from

²⁶*Punch and Judy*, ch. 2 "System Four Is Not a Pair of Eyes", section "Action selection: the function the diagram omits" (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

²⁷Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14-37, p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4. The distinction between tacit knowledge and weak signals follows that note.

the Viable System Model: at any moment an organization could do many things, most of them incompatible, and “something has to decide which one is expressed”.²⁸ The problem, in its words, is “how a field of possible actions is held back, compared, and selectively released”. Its picture of that mechanism is drawn from the anatomy of vertebrate brains, in which the default is inhibition and selection works by releasing one candidate, and the book marks that picture as an analogy, as *Punch and Judy* does.²⁹

The two selection problems sit at the two ends of the loop drawn in chapter 1. Attention is selection under scarcity on the input side; action selection is selection under scarcity on the output side. They fail in related ways. An organization whose attention is captured by the familiar will also find that its actions are the familiar ones, since the model that filters the cues is the same model that proposes the acts. Chapter 6 takes the output side. It matters here only because the design of attention and the design of choice are one design problem seen from two ends.

Machines: pattern detection changes the task of noticing

The obvious modern answer to the filters is a machine that looks at everything. Two things in this chapter’s argument say why that answer is incomplete.

The first is that a machine looks through a partition too. A dashboard encodes the picture that was current when it was built; chapter 1 said that a measurement system is excellent at tracking movement along the axes it expects and poor at reporting that the axes have changed. A model trained on last year’s data has last year’s distinctions. Neither is useless. Both are attention devices with a model inside, and the model was drawn by someone.

The second comes from Hutchins, and it is the more interesting. He argued against the habit of describing

²⁸*Punch and Judy*, ch. 2, section “Action selection: the function the diagram omits”.

²⁹*Punch and Judy*, ch. 2, section “Action selection: the function the diagram omits”.

tools as amplifiers of thought: “None of the component cognitive abilities has been amplified by the use of any of the tools. Rather, each tool presents the task to the user as a different sort of cognitive problem requiring a different set of cognitive abilities”.³⁰ Applied to perceiving, this means that a pattern-detection system changes what attending is; whether it makes the organization more attentive is a separate question. Dave Snowden’s SenseMaker, to take a tool designed for this purpose, collects large numbers of short accounts from people inside a situation and has them index their own stories, so that patterns can be seen across thousands of them that no single reader would see.³¹ What the organization then has to notice is different: the shape in the aggregate, and what that shape does and does not mean, where before it was the single story. The task of noticing has moved. It has not gone away, and the tacit knowledge it requires is of a new kind. Chapter 13 takes up what happens when models of this kind sit inside the loop as participants.

Who may raise a cue, and to whom

Every activity in this book has a collective form, and for perceiving the collective form is a question of permission. Who may raise a cue, and to whom?

Put that way, most of the chapter’s material becomes design. The partition is drawn by someone; it can be drawn by people who see the edges. Channels are set; they can be set so that an unclassified impression can travel without being interpreted first. Silence has causes; they can be removed, or at least named. Hutchins adds a detail that is easy to miss. What a person can learn from the work around them depends on their horizon of observation, and “the design of a tool may change the horizon of observation of those in the vicinity of the tool”.³² A

³⁰Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

³¹Snowden, D. J., SenseMaker (Cognitive Edge / The Cynefin Company).

³²Hutchins (1995), ch. 6, p. 270.

tool, a layout, a meeting format decides what its user sees, and also who else can see it. An open ticket queue and a private one give the same organization two different horizons. “Humans create their cognitive powers by creating the environments in which they exercise those powers”, he writes, and the sentence holds for attention as much as for computation.³³

Episode — The cue five people saw

“Restarting since the building system was connected.” Tobias, a field technician in region South at Halvern, typed the sentence into the free-text field of a service ticket in the spring after the C4 controller generation was launched. Halvern makes heat-pump systems for commercial buildings and sells them through independent installer partners. The unit in front of him was starting and stopping its compressor far more often than it should. According to the building operator it had done so since the installer linked the heat pump to the building-management system, which lowered setpoints at night and sent its overrides through a gateway.

The form would not save without a cause. It offered five: product defect, installation error, operating error, wear, other. Tobias had found nothing wrong with the controller and nothing wrong with the wiring. The trouble appeared only when the two systems were connected, and the form had no category for that. The installer had made the connection, so he picked installation error. Two other technicians in the region made the same choice that spring. On the support desk Jana heard the sentence from building operators on the phone, almost word for word, and filed those tickets the same way.

Tobias said to his team lead that the category did not fit. The team lead told him to stick to the codes. The

³³Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

free-text field was not carried into any report. Quality's root-cause codes mapped one-to-one onto the five categories, and the service dashboard counted tickets per category, per region, per week. What reached anyone above the team lead was therefore a number: installation errors in region South, somewhat up.

There was no longer a place to say it aloud either. The regional depot mornings had ended the year before, and technicians started from home and worked alone. Tobias posted the observation in the technicians' messenger group for region South, where the two colleagues answered that they had seen it too. The group was not an official channel. Nobody from the second line or from Quality read it.

A few of the cases were escalated, and in early summer they reached Milan, a senior second-line specialist with 19 years in the company. He saw the pattern. To raise it he would have had to describe a fault that involved the C4, which was the flagship of the COO, Dr. Reuter, who also chaired the Product Quality Board, the only forum that could act on such a description. Milan did not press it. He resolved each escalation as it came and wrote nothing that joined them.

By midsummer five people at Halvern had seen the same thing: three technicians, a first-line agent and a senior specialist. The dashboard showed installation error, rising.

Return to Beer's page. The parents had the information and not the model; the chapter has argued that organizations are usually in the same position, with two further filters that a family does not have. The organizational reader of that page has an advantage the parents did not: the partitions, channels and silences of an organization are visible, they were made, and they can be remade. What the parents needed was a distinction they did not have, and somebody who had it and was in a position to say so.

The cues that pass all three filters arrive somewhere, in someone's attention, as elements of a situation. Chapter 3 asks what a picture that they fit consists of.

What would count against it

The chapter claims that what an organization notices is set by its model, its channels and its safety to speak, and that these are designed. Against it would count organizations that change their agendas, dashboards, forums and reporting channels and notice nothing different afterwards. Against the claim that detection is a distributed social activity would count central horizon-scanning units that detect change earlier than distributed frontline networks in comparable settings. Where to look: the high-reliability studies of Weick and Sutcliffe, Vaughan's account of normalized deviance, and the research on organizational silence that followed Morrison and Milliken.^{34,35,36}

References

- Ansoff, H. I. (1975). Managing strategic surprise by response to weak signals. *California Management Review*, 18(2), 21–33.
- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Conant, R. C., & Ashby, W. R. (1970). *International Journal of Systems Science*, 1(2), 89–97.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

³⁴Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the Unexpected: Sustained Performance in a Complex World* (3rd ed.). Wiley.

³⁵Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.

³⁶Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725. p. 706 (organizational silence defined), p. 708 (a climate of silence: futile and/or dangerous).

- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Mathis, K. (2025). *The Tautai Principle*.
- Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Pfiffner (2017).
- Polanyi, M. (1966). *The Tacit Dimension*.
- *Punch and Judy*.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059–1074.
- Snowden, D. J., SenseMaker (Cognitive Edge / The Cynefin Company).
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.
- Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the Unexpected: Sustained Performance in a Complex World* (3rd ed.). Wiley.

Chapter 3

Chapter 3 — Knowing what is going on: situation awareness

A theory in three levels

In 1995 the journal *Human Factors* published a paper by Mica Endsley under a careful title: “Toward a theory of situation awareness in dynamic systems”.¹ The care was in the first word. What the paper delivered, and what has been in use since, is a definition in three parts, and chapter 1 gave it in a sentence: perception of the relevant elements of a situation, comprehension of what they mean together, and projection of their state into the near future.

The three parts are a ladder, and each rung stands on the one below. Perception, Endsley’s Level 1, is the arrival of the elements: the reading on the gauge, the number on the report, the remark in the meeting, the customer who did not renew. Comprehension, Level 2, is what those elements mean together, in relation to a goal: that the readings add up to a system running hot, that the numbers and the remark together mean a segment is moving away. Projection, Level 3, is where this is going: what the situation will look like in a few minutes, or a few quarters, if nothing is done, and what it will look like if something is.² A person can have the first without the

¹Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543> p. 36 (definition: perception, comprehension, projection), p. 38 (relevance depends on the operator’s tasks and goals).

²Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64.

second, and the second without the third. That is the point of keeping them apart.

One word in the definition does more work than it seems to. Perception is of the *relevant* elements, and relevance is decided by what the perceiver is trying to do. Endsley's awareness is awareness for a purpose: the pilot's, the controller's, the operator's, and the same instrument reading is relevant to one purpose and noise to another.³ In an organization the purposes are several, and they are held by different units. The elements relevant to the sales unit's goal differ from those relevant to the finance unit's, and the two comprehend the same quarter differently, each correctly for its purpose. This is why the picture across an organization is never one picture, and why the second half of this chapter has to ask how pictures held for different purposes are made to overlap where they must.

What the ladder is for, in practice, is diagnosis. *They did not know what was going on* is the verdict after most organizational failures, and it says nothing that could be acted on. Endsley's levels turn it into a question with three possible answers. Were the elements never perceived, because nobody was looking there, or because the channels described in chapter 2 did not carry them? Were they perceived and not understood, because the model in use had no place for what they meant together? Or were they perceived and understood, and the projection was wrong, so that the organization acted on a future that did not arrive? Each answer points to a different remedy, and the remedies are not interchangeable. An organization that failed at Level 1 does not need a better analyst. One that failed at Level 3 does not need more data.

Read this way, the two cases of chapter 1 fail at different rungs, and the reading is mine rather than Endsley's

<https://doi.org/10.1518/001872095779049543> p. 36 (definition: perception, comprehension, projection), p. 38 (relevance depends on the operator's tasks and goals).

³Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543> p. 36 (definition: perception, comprehension, projection), p. 38 (relevance depends on the operator's tasks and goals).

or Weick's. The crew at Mann Gulch perceived the fire and comprehended it as a fire of a known kind; where the picture failed was projection, in a future in which the fire did what fires of that kind do, and it did not. The navigation team on the *Palau* lost its instrument and kept all three levels, because the levels were not in the instrument. Where they were is the subject of the second half of this chapter.

The demons of awareness are organizational

Endsley and Debra Jones, in the second edition of *Designing for Situation Awareness*, list eight ways in which awareness typically breaks down, and they call them the demons of SA.⁴ The book was written for the designers of cockpits, control rooms and interfaces, and its demons are described at the scale of one operator and one display. Read at the scale of an organization they lose none of their force, and this section takes five of them across, in my own words.⁵

Attentional tunneling is the first. An operator locks on to one part of the situation and stops sampling the rest. The organizational version is a management forum that has spent three months on one topic, an integration, a crisis, a launch, and in those months has not looked at the thing that was moving elsewhere. Nobody decided to stop looking. The agenda filled, and attention followed it, exactly as chapter 2 described.

Data overload is the second, and chapter 2 has said most of what needs saying about it: Simon's poverty of

⁴Endsley & Jones (2012), ch. 3, p. 31: the eight demons of situation awareness, "errant mental models" among them.

⁵The five demons discussed (attentional tunneling, data overload, misplaced salience, errant mental models, out-of-the-loop) are named from the outline and described in the author's words; "errant mental models" is at p. 31, the remaining names and the descriptions of each demon, including the out-of-the-loop syndrome, follow Endsley & Jones (2012), ch. 3.

attention is a Level 1 failure in Endsley's terms, since elements that were available were never perceived because too many others were.⁶

Misplaced salience is the third. Displays draw the eye to what is bright, red or moving, whether or not that is what matters, and the eye goes. The organizational version is the dashboard whose one red number captures the meeting, and the risk register whose entries are ranked by the ease with which they can be scored rather than by what they would do. The salient element is perceived; the elements that would have changed comprehension were dull, and were not.

Errant mental models are the fourth, and this one the book has been circling since chapter 1. A person who holds the wrong model perceives the elements and comprehends them wrongly, and the wrong comprehension produces a wrong projection with the same confidence as a right one. This is Beer's parents from chapter 2, seen from the other side: they were not short of Level 1. They had no Level 2 for it, and so their Level 3, whatever they expected of the coming months, was built on a situation that was not the one they were in.

Out-of-the-loop is the fifth. When a task has been automated, the person who is supposed to supervise it loses the awareness that doing the task used to produce, and when the automation hands back control, the person has to rebuild the picture from nothing at the worst moment.⁷ Chapter 13 takes this up in its own terms, because a model that sits inside the loop and does part of the perceiving and comprehending produces exactly this condition for

⁶Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press. The poverty-of-attention sentence and the design rule are in the OCR'd scan (printed pp. ~40–42); see chapter 1.

⁷The five demons discussed (attentional tunneling, data overload, misplaced salience, errant mental models, out-of-the-loop) are named from the outline and described in the author's words; "errant mental models" is at p. 31, the remaining names and the descriptions of each demon, including the out-of-the-loop syndrome, follow Endsley & Jones (2012), ch. 3.

the people around it. The organizational form is older than automation, though. A manager who receives only a summary is out of the loop with respect to everything the summary compressed, and it is when the summary stops fitting that the manager discovers how much that was.

The value of the list for an organization designer is that each demon names a mechanism, and mechanisms can be looked for. The demons are not personal failings. They are properties of the arrangement of people, tasks and displays, and the arrangement can be changed.

They also have a collective form, and it is worse than the individual one. Hutchins, in the study the book returns to below, asked what an individual property such as confirmation bias produces when several people who share it communicate, and his simulations showed that the answer depends on how the communication is organized, with richer communication under some conditions producing worse properties at the level of the group.⁸ A shared errant model is of this kind. One operator with the wrong model can be corrected by another with the right one; a forum in which everyone holds the same wrong model corrects nobody, and each member's confidence is raised by the agreement of the rest. The organizational demon is a group of people, arranged so as to confirm each other, who misread the situation together and take the agreement for evidence.

Mental models as the enablers of comprehension and projection

Endsley and Jones make one statement about the three levels that this book leans on more than any other. Mental models, they write, are "the key enablers of Levels 2 and 3 SA".⁹ Perception delivers elements. What turns elements into meaning, and meaning into expectation, is a model of how the situation works.

⁸Hutchins (1995), ch. 5, pp. 239 and 252.

⁹Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press, pp. 21–22.

This is where the internal model of chapter 1 meets situation awareness, and the meeting is exact. In Endsley's vocabulary a model is what the operator holds of the system being operated and the environment it operates in; in this book's vocabulary it is the account an organization acts on, of its environment and of itself, held in heads, artifacts and structure. The claim is the same at both scales. Without a model there is Level 1 and nothing above it. With a poor model there is confident comprehension of the wrong situation. With a good model, comprehension is fast and projection is possible, and people who hold such models look, from outside, as if they simply see what is going on.

Several authors the book uses describe the forward use of the model in almost the same words, from very different traditions. Gary Klein, studying fireground commanders and other experts, found them using what he calls mental simulation, "the ability to imagine people and objects consciously and to transform those people and objects through several transitions", running a situation forward to see what a course of action would produce, and backward to explain something strange.¹⁰ Stafford Beer, writing about the part of an organization that looks ahead, calls foresight "rapid-time simulation" and insists that the unit responsible for it "must contain a model — some model — of the corporation".¹¹ Andy Clark, summarizing a research program in cognitive science, describes brains as "prediction machines" running a "generative model" of their inputs.¹² The book carries Clark over to organizations as a comparison only, and says so wherever it does. But the convergence is real: projection, the top rung of Endsley's ladder, is a model run forward, and it is available only to the extent that the model is good.

The dependency runs the other way too, and Endsley's own framing makes it visible. A model decides which

¹⁰Klein (1998), ch. 5 "The Power of Mental Simulation" (from p. 49).

¹¹Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

¹²Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, pp. 3–4.

elements are relevant, and therefore which are perceived. The categories of Level 1 are supplied by the model that will later be used at Level 2. This is Ashby's observer from chapter 2 in another vocabulary, and it explains why errant mental models are the most dangerous of the demons: they corrupt every rung at once, and from the inside there is no symptom.

For a practitioner the useful consequence is that models can be inspected through projections. A model is mostly tacit and hard to state, as chapter 2 said of the knowledge that lets a person recognize a signal. A projection is easier to state: ask the people who hold a picture of a market, a program or a site what they expect to see in three months, and the differences in their answers expose differences in the models they hold that no amount of discussion of the current situation would have brought out. Where the projections agree and the arrivals differ from them, the shared model is wrong somewhere, and the place is visible. Where the projections disagree, the organization has learned that it holds more than one model and can ask which distinctions they differ on. Neither of these is available to an organization that never states its projections, and most do not, because a stated projection can be seen to fail.

The book is aware of a dissent here that it does not resolve in this chapter. Weick, whose account of sense-making chapter 4 sets out, is suspicious of models and trusts action; his soldiers found their way with the wrong map. Chapter 5 says where the book stands between that view and the one taken here. In this chapter the position is only the one Endsley and Jones state: without a model there is no comprehension and no projection, and awareness stops at Level 1.

Team, shared and distributed awareness

Endsley did not stop at the individual operator. Her later work extends the model to teams, and it does so in a way that preserves the individual as the place where awareness sits. Team situation awareness, in her account,

is the degree to which each member holds the awareness that their own role requires; shared situation awareness is the overlap, the part of the picture that two members both need and both must hold in the same way.¹³ Both are useful in an organization. The first asks whether the person who has to act on a segment shift has the picture that acting requires. The second asks whether sales and product, or the site and the center, hold the same picture where their work meets, and it explains a familiar kind of failure: two units each with an adequate picture that do not agree about the part they share.

A second line of research goes further, and it grows out of Edwin Hutchins' work. If, as Hutchins argued, the unit that thinks is the system of people and artifacts rather than the person alone, then awareness is a property of that system, and it can be held in it even where no member holds it whole. Hutchins moved "the boundaries of the cognitive unit of analysis out beyond the skin of the individual person" and treated "the navigation team as a cognitive and computational system".¹⁴ The research on distributed situation awareness takes that step for the concept this chapter is about; the volume that McNeese, Salas and Endsley edited in 2021 is where the two lines meet, and Endsley's name among its editors says that the two lines are in conversation.¹⁵ In it Orestis Palermos states the general claim of the field: appropriately interacting individuals can form "collective entities with cognitive properties, over and above the sum of their individual members' cognitive properties".¹⁶ And Jennifer Reep, reviewing the

¹³Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64, p. 39 ("the degree to which every team member possesses the SA required for his or her responsibilities"); Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press, ch. 11, pp. 195–196 (team and shared SA).

¹⁴Hutchins (1995), introduction (chapter summaries; the local functional system).

¹⁵McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.

¹⁶Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 5, p. 115). CRC Press.

expertise literature in the same volume, adds that experts distribute their cognition into their context and artifacts in ways novices cannot yet manage, which means that the picture an expert holds is partly held outside the expert.¹⁷

Hutchins gives the picture an image that this chapter will keep. Because the members of the navigation team knew overlapping parts of each other's work, the team could "reconfigure dynamically" when a part failed, and "the individuals are a sort of flexible tissue that moves to ensure the propagation of task-relevant representational state".¹⁸ Overlapping knowledge, he says, is "an important source of the robustness of such systems in the face of error and interruption". Awareness, on this reading, is held by a tissue, and the tissue moves. When the gyro-compass on the *Palau* stopped, no single head held the ship's position, and the position did not disappear; the tissue closed around the gap and kept the picture, minute by minute, until the anchor went down. Chapter 1 told that story. Here it is the demonstration that Level 1, 2 and 3 can survive the loss of the instrument that supplied Level 1, because they were never in the instrument alone.

Here the chapter marks a disagreement it does not settle. Endsley places awareness in an operator's head, with displays, tools and other people as inputs to it. Hutchins places it in the functional system and treats the head as one of its media. The two are not simply compatible, since they answer differently when asked where the picture is when nobody holds it, and each has a practical consequence. The Endsley reading directs design at the individual's inputs: what does this person need to see? The Hutchins reading directs design at the system: where is each part of the picture held, which handoff carries it, and what fails when a part is removed? This book uses both and says which it is using. Part III makes the Hutchins reading

¹⁷Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 8, pp. 187–191). CRC Press.

¹⁸Hutchins (1995), ch. 4, pp. 221 and 227.

its core, because organizations are the case in which nobody holds the picture whole. For the individual operator, Endsley's account remains the more precise.

The distributed reading turns a vague worry into a set of questions that can be asked of any unit. Which parts of its picture are held in heads, and whose? Which are held in artifacts: a board, a log, a shared document, a dashboard, a chart? Which are held in the structure itself, in the fact that a certain role exists and reports to a certain other? And at which handoffs does a part of the picture have to move from one medium to another, so that it can be lost there? The last question has a temporal form that every organization with shifts, rotations or project handovers knows: awareness that was held by a tissue of people at four o'clock has to be held by a different tissue at five, and the artifacts that carry it across, the log, the handover note, the board, are the whole of the picture for the minutes in which nobody holds it. Chapter 7 follows a single piece of the picture across such handoffs. Here the claim is only that awareness in an organization is held that way, in a system, and that the system can be looked at.

There is also a limit to what Hutchins brings, and it should be stated here rather than discovered later. His material is well-structured work in which the situation is not in question. Where the ship is, is a matter of computation; nobody on the bridge doubted that there was a position to be found. Ambiguity, surprise and the collapse of meaning are not in his study, and they are the subject of the next chapter.

Projection: the prospective side of understanding

The top rung of the ladder deserves a section of its own, because it carries the book's temporality thread and because organizations treat it in two contradictory ways.

Projection is short-range, continuous and revisable. In Endsley's operators it reaches a few seconds or minutes ahead, it is updated as new elements arrive, and it is held lightly: an expectation, not a commitment. The organizational activity that most resembles it is the running

expectation that experienced people carry of what the next weeks will bring, and revise without ceremony when they bring something else; the strategic plan resembles it much less. That expectation is the model run forward, and it is what lets people notice a surprise as a surprise: the element that arrives is compared with the element that was projected, and the difference is the signal. Without projection there is no surprise, only a sequence of facts.

The contradictory ways in which organizations treat it are these. On one side, projection is neglected as a skill and left to whoever happens to hold a good model, without a forum in which projections are stated and compared. On the other side, it is formalized into forecasts and plans that are then held as commitments, so that a projection that turned out wrong becomes a failure to be explained rather than a cue to be used. Chapter 1 quoted Beer's remark that a corporate plan "must continuously abort" if the match between the world and the firm's notion of it is to be kept.¹⁹ A projection is a plan that expects to abort. It is made in order to be compared with what happens, and its value lies in the comparison. Chapter 6 gives this a further turn, following Boyd, in which a decision itself is read as a hypothesis and the action as its test; a projection stated in advance is that hypothesis in its mildest form.

Two other chapters take the forward use of the model further. Chapter 5 develops orientation as the continuous revision of the model in both directions, forward as projection and planning, backward as explanation, and settles where the book stands on what a model is for. Chapter 4 supplies the backward direction in its strongest form, since Karl Weick's sensemaking is retrospective by construction: it makes sense of what has already happened. Projection is the prospective half of understanding, and this chapter's claim is only that it exists, that it depends on the model, and that an organization can be asked where it does its projecting and whether anyone compares the projections with what arrived.

¹⁹Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

The limit: a situation that exists, and categories that are available

Situation awareness assumes two things, and the assumptions mark where the lens stops.

It assumes, first, that there is a situation to be aware of. The elements are there; the task is to perceive them, comprehend them and project them. That is true of a cockpit and of a ship's bridge, and it is true of most of an organization's working days. It is not true of the occasions chapter 1 listed under surprise, equivocality and novelty, where the elements support several readings at once, or belong to no category the organization has. On those occasions there is no picture to be accurate about until one has been made. The organizational cases are familiar enough: a competitor of a kind that fits none of the segments in the market model, a regulation that cuts across the categories of the risk register, a merger partner whose way of working matches no role in the organization chart. In each case the elements arrive, and Level 1 is satisfied. What is missing is a Level 2 in which they could mean anything in particular, because the model that would supply the meaning has not been built.

It assumes, second, that the categories for perceiving are available. Level 1 is perception of the *relevant* elements, and relevance is decided by a model that already exists. Where the model has no category, there is no element, only noise, and the account chapter 2 gave of the invisible signal applies. Endsley's operator with an errant mental model is still an operator with a model. The harder case is the one with no model for what is happening, and that case is not a demon. It is the end of the ladder.

Stress narrows all three levels, and this is where the book's thread on emotion first touches the argument. Endsley and Jones treat stress, anxiety and related pressures among the conditions that degrade awareness, at perception, at comprehension and at projection alike.²⁰ Under

²⁰Endsley & Jones (2012), on stress, anxiety and related pressures as conditions that degrade SA. The claim that stress narrows all three levels is from the outline.

threat, attention tunnels, models are held more rigidly, and the horizon of projection shortens to the next few moments. The organizational form is a unit under pressure that perceives less, understands what it perceives in the terms it already has, and projects only as far as the end of the quarter. Chapter 1 quoted Weick's name for the extreme case, the cosmology episode, in which people "suddenly and deeply feel that the universe is no longer a rational, orderly system".²¹ Read on Endsley's ladder, a cosmology episode is the failure of all three levels at once, with the model that supplied them gone. That is where awareness ends and something else has to begin.

Episode — Perceived, not comprehended

In one cold week in January the service dashboard at Halvern, a manufacturer of heat-pump systems for commercial buildings, showed 14 failed compressors in six days. Next to that number stood the column for *installation error*, which had been climbing since the summer. Henrik, the Head of Quality, read the two together: installer errors, made visible by the cold. The service leadership team took that account the same morning, and the question in the room shortened to the next cold week. The proposal that followed was for more data. A new monitoring dashboard would show compressor starts per unit and per day, so that the next failure could be seen coming.

Halvern had already acted on a future. In November it had launched a training program for its installer partners, at about €180,000, on the expectation that installation errors would fall once the partners were trained. Parts planning had stocked compressors for a normal winter, about 25. By the end of the season 96 had been replaced.

In March Katrin, the Head of Service, prepared three traced cases for the Product Quality Board, and for

²¹Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

each she asked what had been known and when. The answer was the same three times. In the spring Tobias, a field technician in region South, had written into the free-text field of a ticket that the unit had been “restarting since the building system was connected”. Jana had heard that sentence from building operators on the support desk. The tickets had been filed, counted and shown on the dashboard every week. Nothing had been missing at Level 1. What Halvern lacked was a model in which a controller and a building-management system, each doing what it was built to do, could wear out a compressor between them. The ticket form offered five cause categories, the installer had connected the building system, and so each of these tickets became an *installation error*. Comprehension followed the category, and the projection for the winter was run forward from it with the usual confidence.

The diagnosis changed the remedy. The monitoring dashboard was put on hold: it would have delivered start counts to people who would have sorted them under the same five categories. The work went into the account instead. Milan, a senior specialist on the second line, set out for the Board how the building system’s setpoint overrides made the new controller start and stop its compressor far too often, and the Board replaced the installer account with that of an interaction fault. What followed was a longer minimum off-time, a changed gateway setting and, in June, a firmware fix. None of it needed a reading that Halvern did not already have.

For most of a year the technicians who had the cue had no category to put it in. What they made of it went into the free-text field and into their messenger group.

The tissue can move only toward something. On the *Palau* it moved toward a position that everyone knew existed and that the team knew how to compute; the categories were intact and the loss was of an instrument. At

Mann Gulch there was no longer a category for what the fire was doing, and the crew, which had been a tissue of the same kind, had nothing to move toward and came apart. Situation awareness describes the first case completely and the second not at all. Where cues fit no available category, awareness runs out, and the people who are left with the cues have to make something of them. That activity has a name of its own, Weick's, and it is the subject of chapter 4.

What would count against it

The chapter claims that the three levels locate where awareness fails, and that the location changes the remedy. Against it would count organizational failures of awareness that cannot be located at a level, so that the three-level model adds nothing to the diagnosis beyond the verdict that people did not know. Against the claim that mental models enable the upper levels would count comprehension and projection that are found to be as good in people without a model of the situation as in people with one. Against the distributed reading would count teams that lose a member or an instrument and lose the picture with it, in settings where the picture was supposedly held by the system. Where to look: Endsley's own empirical work on awareness errors, and the aviation and control-room studies of breakdowns in awareness.^{22,23}

References

- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.

²²Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543> p. 36 (definition: perception, comprehension, projection), p. 38 (relevance depends on the operator's tasks and goals).

²³McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.

- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543>
- Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.
- Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In McNeese, Salas & Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*, ch. 5. CRC Press.
- Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In McNeese, Salas & Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*, ch. 8. CRC Press.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press.
- Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.

Chapter 4

Chapter 4 — When the picture breaks: sensemaking

A map of the wrong mountains

A Hungarian reconnaissance unit is lost in the Alps. It has a map, and with the map it finds its way back.

The map is of the Pyrenees.

Karl Weick tells the story on pages 54 and 55 of *Sensemaking in Organizations*, and his explanation of why the unit came back is exact about what the map did and did not do. The soldiers succeeded because they were active, because they had a purpose, and because “they had an image of where they were and where they were going. They kept moving, they kept noticing cues, and they kept updating their sense of where they were.”¹ His conclusion is the line chapter 1 quoted: “when you are lost, any old map will do”.²

Put the explanation beside the situation it explains. The map did not show the mountains the unit was in. What it gave the unit was a reason to move in one direction rather than to wait, and a frame in which whatever the soldiers then saw could be read as confirmation, or as a cue that the frame needed adjusting. The wrong map worked because the unit did not treat it as the terrain. It treated it as a first account, and it kept revising.

Being lost is where Weick’s theory begins. Chapter 3 described situation awareness, Mica Endsley’s account of holding a picture that fits: perceiving the elements, comprehending what they mean together, projecting where

¹Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 54–55.

²Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 54–55.

they lead.³ Awareness assumes that a situation is there to be perceived and that the categories for perceiving it are at hand. The unit in the Alps had a map that did not describe the ground it stood on, and no better one. What people do then, when the picture they hold does not fit and no ready picture replaces it, is what Weick calls *sensemaking*. The book uses the word in his sense only, and for the rest of the loop it says *understanding*.

Where awareness runs out

Chapter 1 named three occasions that call for sensemaking: surprise, when something happens that should not have; equivocality, when the cues support several readings at once; novelty, when the situation belongs to no category the organization has. Weick has a word for the way all three are met from inside: interruption. An ongoing activity is cut into by something it did not expect, and the flow that carried the activity, and the account that went with it, stops.⁴ The smokejumpers at Mann Gulch were interrupted, and so was the navigation team on the *Palau* when the gyrocompass failed. One crew rebuilt a picture and the other did not, and chapter 1 said why.

The distinction that matters most in practice is the one between equivocality and uncertainty, because organizations habitually treat the first as if it were the second. Under uncertainty, information is missing and can in principle be gathered; the remedy is more data, and the analyst's question is where to get it. Under equivocality the information is there, and it supports two or more readings. More data does not help, because the problem is too many meanings, and what is needed is a way of settling on one that is good enough to act on.⁵ A drop in a satisfaction score, a competitor's odd hiring pattern, a customer

³Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. Chapter 3 gives the model in full.

⁴Weick (1995), on interruption as the occasion for sensemaking.

⁵Weick (1995), equivocality as distinct from uncertainty; p. 27 (equivocality distinct from uncertainty).

who starts using new words for an old complaint: none of these is a missing fact. Each is a fact with several plausible meanings, and the meeting that treats it as a data problem commissions a report and waits.

There is a simple test for which of the two one is in. Give the same report to two competent people who share the organization's categories. If they disagree about what is in it, the problem is a missing fact and the remedy is to find it. If they agree about what is in it and disagree about what it means, and neither can name a further fact that would settle the matter, the situation is equivocal, and a further report will be read in the same two ways. The test is not exact, and it is enough to stop a management team from spending a quarter on data collection when what it needed was an hour of argument.

Ikujiro Nonaka, whose theory of knowledge creation the book uses in Part IV, starts from the same place under another name. He speaks of "fluctuation" in the environment and of "breakdowns", borrowing the term from Winograd and Flores, as the occasion on which individuals try to "make sense of their environment" and, in doing so, create knowledge.⁶ For Weick the interruption occasions sensemaking; for Nonaka it occasions knowledge creation. The two theories part company on the standard the account has to meet, and the chapter comes back to that.

Seven properties, and what each is for

Weick's 1995 book organizes sensemaking around seven properties. They are paraphrased here.⁷ Ten years later Weick, Kathleen Sutcliffe and David Obstfeld restated

⁶Nonaka (1994), p. 17 (individuals create knowledge; the organization amplifies it; intention, autonomy and fluctuation; Weick 1979 on enactment), p. 18 (breakdowns after Winograd and Flores). The gloss of "breakdown" as the moment a routine stops working and its assumptions become visible is the author's; Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

⁷Weick (1995), ch. 2, the seven properties (identity, retrospect, enactment, social, ongoing, extracted cues, plausibility). Paraphrased.

the theory with organizing in the title.⁸ The properties are a description, assembled from many sources, and the reason to go through them one by one is practical: each names a way in which a familiar organizational habit misreads what people are doing when the picture breaks.

Sensemaking is *grounded in identity*. Who I take myself to be decides what I can make of a situation, and a situation that threatens the identity threatens the capacity to make sense of it. This is the property Mann Gulch turns on. The order to drop the tools, chapter 1 argued, was for many of the men an order to stop being firefighters, and once the role structure went, the frame went with it. In organizations the identity in question is usually collective. A unit that has understood itself as the market leader reads a competitor's move differently from one that has understood itself as the challenger, and a reorganization that changes what a unit is changes what it can understand before it changes anything else. Chapter 5 gives the property a structural place: an organization has to hold a model of its environment and a model of itself, and Stafford Beer, whose Viable System Model the book uses in Part III, requires the second as firmly as the first.⁹

Sensemaking is *retrospective*. People know what they have done, and what it meant, after they have done it. The sharpest form of the claim is Weick's reading of Harold Garfinkel's study of jurors: "the person defines retrospectively the decisions that have been made".¹⁰ What looks in the minutes like a decision followed by its execution was often an action followed by an account of why it had been the decision. The property explains why the accounts given in reviews are so much tidier than the events they review, and chapter 6 draws the organizational consequence: without a record made at the time, the organization rewrites

⁸Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization Science*, 16(4), 409–421; p. 409 ("sensemaking is about the interplay of action and interpretation").

⁹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

¹⁰Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 10–11, on Garfinkel's jurors.

its choices after the fact, and learning has nothing to test against.

Sensemaking is *enactive*. Chapter 1 took this property as one of the book's premises: organizations help produce the environment they then try to read. Here it has a practical edge. Waiting is also an action, and the situation that is eventually analyzed is partly the product of the wait.

Sensemaking is *social*. An account is built in talk, with others, and with an audience in mind even when one is alone. The unit of sensemaking is therefore the conversation and the meeting, and the design of those, who is in the room, who speaks first, what may be said, is the design of understanding. Part III is about that.

Sensemaking is *ongoing*. It has no clean start. People are always in the middle of something, and an interruption cuts into a flow that was already carrying its own account of what was going on. An organization that schedules sensemaking as a phase, after the analysis and before the plan, has mistaken an activity for a stage. The property also sets the book's treatment of time. Sensemaking looks backward by nature, the prospective sensemaking of a later section looks forward, and chapter 5 asks at what rhythm an organization returns to its account; the three belong together, and a planning calendar that fixes the third without providing for the first two has fixed the least important of them.

Sensemaking is *focused on and by extracted cues*. A cue is a small element of a situation that is noticed, pulled out, and made to stand for a larger pattern. Which cues get extracted depends on the frame and on the identity, which is why two departments can look at the same event and extract different cues from it, and why chapter 2 treats attention as the first design problem of the loop.

Sensemaking is *driven by plausibility rather than accuracy*. The account that gets people moving is the one that is coherent enough and shared enough to act on, and its accuracy is settled later, if at all, by what the action produces. The Hungarian unit's map was plausible. It was not accurate, and it did not need to be.

Read as a set, the seven describe an activity that is continuous, social, retrospective and satisfied with plausibility. Each of these is the reverse of what a planning cycle assumes: that understanding happens first, in a phase, by analysis, and is finished when it is accurate. The planning cycle is a good instrument where its assumptions hold. It is poorly placed to notice when they have stopped holding, because noticing that is sensemaking, and the cycle has no stage for it.

Not analysis

Two things sensemaking is regularly confused with need to be set apart from it, because the confusion leads to the wrong remedy.

The first is analysis. Analysis takes a situation that can be broken into parts, studies the parts with data, and re-assembles a conclusion. It assumes that the situation stays still while it is studied and that the categories used to break it up are the right ones. Where that holds, analysis is the right tool and sensemaking is waste; chapter 1 drew the line, with Dave Snowden's Cynefin, between the complicated and the complex. Where it does not hold, analysis produces an accurate answer to a question the situation is no longer asking. Sensemaking is what people do with the part analysis cannot settle: which question this is, which categories apply, what the numbers are numbers of.

The second is interpretation, in the sense of reading a meaning off data that is already given. Chapter 1 argued that data does not say what it means. Weick's point goes further. In sensemaking the data are not given first and read afterwards; the frame decides which cues are extracted, and the cues then bear on the frame. The two are produced together, and a change in either changes the other.

Gary Klein and his colleagues built a model around exactly this relation and called it the data-frame model. A frame is a structure (a story, a map, a script, a plan) that says which data matter and how they connect; the data, once attended to, question the frame, elaborate it, or force

its replacement.¹¹ Read beside Weick, and this reading is mine, the data-frame model is the tidier of the two: it stays with a frame that is there, however minimal, and follows it as it is kept, elaborated and, when necessary, replaced. Its authors allow for the harder case too — a frame to be built where none exists — so the two accounts differ more in where they begin than in what they permit. Weick begins with the frame's collapse. Both are needed, and the difference between them is roughly the difference between a picture that is being corrected and a picture that has broken.

The frame that breaks

What breaks, when the picture breaks, is the frame; and the frame, in the terms chapter 1 introduced, is the internal model in its loose sense: the account a person or an organization acts on, of the situation and of itself. Sensemaking is what repairs it. The repair has two features that a purely cognitive vocabulary misses, and both are visible at Mann Gulch.

The first is that the collapse is felt before it is understood. Weick's name for the extreme case is the cosmology episode, in which people "suddenly and deeply feel that the universe is no longer a rational, orderly system".¹² The verb is *feel*. Sally Maitlis, Timothy Vogus and Thomas Lawrence have reviewed the place of emotion in sensemaking and argue that emotion belongs to the process rather than disturbing it: it can start sensemaking, and it shapes what account can be built.¹³ The practical form of

¹¹Klein, G., Moon, B., & Hoffman, R. R. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, 21(5), 88–92; p. 88 (frames as stories, maps, organizational diagrams or scripts; elaborating, questioning, re-framing). The comparison with Weick is the author's reading.

¹²Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

¹³Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247. p. 222: emotion "signals the need for and provides the energy that fuels

the claim is that fear narrows. A frightened group extracts fewer cues, entertains fewer readings, and holds on to the identity it has rather than building the one the situation needs. The order to drop the tools was heard by frightened men, and the escape fire was seen by them. Any forum an organization designs for making sense under pressure has to be designed for people in that state, which is a different design task from designing for analysts at a desk.

The second is that the frame includes the role structure. Weick's analysis of Mann Gulch ties the collapse of sense to the collapse of the crew's structure; groups put together for one job are "susceptible to sudden losses of meaning" because there is nothing left to hold a picture in once the frame and the roles go together.¹⁴ The organizational reading is direct. When a reorganization removes the roles through which a unit knew what it was, it removes part of the unit's capacity to make sense, at the moment it asks that unit to make sense of the reorganization.

Now the map. Weick's story separates two things a frame does, and the separation is one the rest of the book needs. A frame gets people moving: it makes one direction more plausible than none, and it makes the next thing seen readable. Call that the *orienting* model. A frame also fits, or fails to fit, the terrain, and the fit is what makes steering by it better than steering by nothing. Call that the *regulating* model. The unit in the Alps had the first and, at the start, not the second, and the second was built on the move, out of cues, by a unit that kept updating. Weick's suspicion of models, which chapter 1 noted, is a suspicion of the second when it is demanded first: of the requirement that the picture be accurate before anyone moves. The cybernetic case for models, which chapter 5 states, is a case for the second: a regulator steers no better than its model of what it steers. Both are right about their

sensemaking" and affects the plausibility of accounts; p. 227: intense emotions "may forestall sensemaking...", moderate emotions energize it most.

¹⁴Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

own half. The map got the unit walking, the updating got it home, and Beer's remark that a corporate plan must continuously abort, which chapter 1 quoted, is the bridge between the two, since a plan that aborts is a model that has noticed it does not fit.¹⁵

Chapter 1 introduced two senses in which the book speaks of a model, and the map shows both at work. In the loose sense, the unit had a model from the moment it opened the map: a wrong one, held in a piece of paper and a few heads, and revised as it went. In the strong sense, a regulator needs a model of what it regulates, and the unit became such a regulator only as the map was corrected against the mountains. Chapter 5 states the two senses in full, gives the warrant for the strong one as far as it reaches, and says where the book uses which. What this chapter adds is the order of events. The orienting model comes first, and it is allowed to be wrong. The regulating model is what the loop produces if the orienting one is held loosely enough to abort.

Holding a model loosely sounds like a temperament. In an organization it is a set of arrangements, and they can be listed. The account is written down with the assumptions it rests on, so that a mismatch has something to be a mismatch with. Someone is charged with noticing the mismatch, and has a place to say so where saying so is not read as disloyalty to the plan. The account is revisited at a rhythm that is shorter than the time it takes the terrain to change, which chapter 5 treats as a design variable. And the people who hold the account are the ones who see the cues, or are within one conversation of them, which is Part III's subject. The *Palau's* navigation team, chapter 1's counter-case to Mann Gulch, had most of this by training and by the structure of its work, and it rebuilt its picture in the middle of a harbour approach. The smokejumpers had none of it, and the terrain changed faster than the account could.

¹⁵Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

Action as a probe

“Action shapes cognition.” The sentence is Weick’s, from page 12 of the 1995 book, and it reverses the order most organizations assume.¹⁶ The assumption is that one understands, then acts. Weick’s claim is that acting is one of the ways one comes to understand: the action produces consequences, the consequences are cues, and the account is built from them afterwards.

The soldiers’ walking was such an action. It did not execute a plan; it tested the map, and each thing that did or did not appear where the map put it was a result. Chapter 6 gives this a mechanism and a vocabulary. In John Boyd’s full loop, as Frans Osinga reconstructs it, decision is hypothesis and action is test, and decisions “feed back into the systems as validity checks on the correctness and adequacy of the existing orientation patterns”.¹⁷ Andy Clark, from the side of cognitive science, separates actions that change the world toward a goal from “cheap epistemic action” that changes what the agent knows.¹⁸ Weick, Boyd and Clark say the same thing in different vocabularies: every action has two values, what it achieves and what it reveals, and in a broken frame the second is the one that matters.

Garfinkel’s jurors return here with a sharper point. If the decision is defined after the action, then an organization that wants to learn from its actions has to fix the account at the time, before retrospect tidies it. Otherwise every probe is a success in the version told afterwards. This is the reason chapter 6 separates action selection, the continuous and mostly tacit choosing inside the loop, from decision, the explicit and accountable commitment that leaves a record. A record is a probe’s hypothesis written down before the result is in.

¹⁶Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, p. 12.

¹⁷Osinga (2007), ch. 6, fig. 6.1 “The real OODA loop” and the surrounding text.

¹⁸Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, p. 183 (“cheap epistemic action”; “confirm or disconfirm the hypothesis”).

Mann Gulch has a last thing to say about probes. Wagner Dodge's escape fire was an action of exactly this kind: invented on the spot, outside the picture the crew held, and its meaning had to be constructed by men whose frame had just gone. In the book's reading, which is not Weick's own, the fire failed as a probe because there was no shared account left in which its result could be read.¹⁹ An action can only test a frame for people who still have one.

Making sense forward

Retrospect is a property of sensemaking. It is not a limit on what sensemaking can be about. Gephart, Topal and Zhang set out a *future-oriented* sensemaking, in which people construct, in conversation, images of future objects and phenomena and act toward them, and Stigliani and Ravasi studied how a group does that work together.²⁰ The

¹⁹ Author's reading, not Weick's; developed from the Hutchins reading in chapter 1. The facts of the escape fire are as in chapter 1.

²⁰ Gephart, R. P., Topal, C., & Zhang, Z. (2010). Future-oriented sensemaking: Temporalities and institutional legitimation. In T. Hernes & S. Maitlis (Eds.), *Process, Sensemaking, and Organizing*. Oxford University Press. p. 285: "Future-oriented sensemaking is thus sensemaking that seeks to construct intersubjective meanings, images, and schemes in conversation where these meanings and interpretations create or project images of future objects and phenomena." The chapter's "plausible account of what may happen in order to act toward it" is a paraphrase, not their wording. Stigliani, I., & Ravasi, D. (2012). Organizing thoughts and connecting brains: Material practices and the transition from individual to group-level prospective sensemaking. *Academy of Management Journal*, 55(5), 1232–1259. p. 1232: an ethnographic study of a design consulting firm, the transition from individual to group-level sensemaking (group-level confirmed). p. 1234: "referred to as 'prospective' (Gioia, 1986) or 'future-oriented' sensemaking (Gephart, Topal, & Zhang, 2010)" — so the term *prospective* is Gioia's, not Gephart's. Gioia (1986): "Symbols, scripts, and sensemaking: Creating meaning in the organizational experience," in H. P. Sims, Jr. & D. A. Gioia (Eds.), *The Thinking Organization* (Jossey-Bass), pp. 49–74. The term is in the chapter under the section heading "Prospective Sensemaking," pp. 63–64 (index: "prospective, 63–64"). p. 63: "Use of a script thus has a strong prospective element to it. For that reason, script-based sensemaking and action might be usefully distinguished as prospective sensemaking." Gioia's sense is the script-based *special case* of retrospective sensemaking — meaning "prospectively affirmed" (p. 64) — narrower than the forward construction of images of future objects that Stigliani and Ravasi and Gephart and colleagues describe.

term *prospective sensemaking* is Gioia's (1986), taken up by Stigliani and Ravasi; Gephart and colleagues call the same forward orientation future-oriented. The account of the future is made the same way as the account of the past: from extracted cues, in talk, on plausibility, and it is tested by what happens next.

Strategy work is prospective sensemaking most of the time, and calling it that changes what one expects of it. A strategy is a plausible story about a future in which the organization has a place, good enough to commit resources to, and its accuracy will be settled by the actions it licenses. Read this way, the strategy process that demands an accurate picture before anyone moves is asking of a prospective account what the Hungarian unit never asked of its map, and asking for something the account cannot deliver in any case, since the future it describes will be partly produced by the actions the strategy sets off. That is the enactment property, run forward.

Chapter 3 put projection, Endsley's third level, at the forward edge of awareness. Chapter 5 shows projection, prospective sensemaking and Klein's mental simulation as one model run forward. Here it is enough to say that the forward account is built by the same social, plausible, cue-driven process as the backward one, and fails in the same ways. A group that cannot make sense of what has happened is not going to make better sense of what will.

Stories are how accounts travel

An account, once built, is a story. It has actors, a sequence and a point, and it can be told again. The story is the form in which an account leaves the room it was made in, and therefore the form in which it reaches the people who were not there. David Boje's study of storytelling in a firm and Barbara Czarniawska's reading of organizations as narratives are the two references the book relies on for the claim.²¹ Nonaka approaches the same fact from

²¹Boje, D. M. (1991). The storytelling organization: A study of story performance in an office-supply firm. *Administrative Science Quarterly*,

the side of knowledge. Metaphor, he writes, “plays an important role in the externalization process”, the step in which what someone knows tacitly is first put into words for others, and a story is the most common vehicle of that step.²²

For the design of organizations this has two consequences, which Part III takes up. Forums are where stories are compared. A forum that admits one story and treats the others as noise has settled its sensemaking before it started, and chapter 8 asks what a forum has to look like so that frames can compete before anyone commits. And whose story wins is a question of power. Leaders give sense as much as they make it, and chapter 10 is about that.

The collective form runs through the whole chapter, so it is stated once, plainly. Sensemaking is social from the start. There is no individual sensemaking that is later shared; the identity in the first property is a social identity, the cues are extracted in conversation, and plausibility is plausibility to someone. What an organization designs, when it designs forums, reporting lines, meeting formats and the safety to say what one has seen, is the social process in which its accounts get made.

What would count against it

The chapter claims that in equivocal situations a plausible account, acted on and corrected, serves an organization better than a demand for accuracy before action. What would count against it is evidence that accuracy-first analysis consistently outperforms plausibility-and-action cycles in situations where the equivocality is established rather than assumed: where it has been shown that the

36(1), 106–126. p. 106: stories are “terse, requiring the hearer to fill in silently major chunks”; “organization as a collective storytelling system”. Czarniawska, B. (1997). *Narrating the Organization: Dramas of Institutional Identity*. University of Chicago Press.

²²Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

cues supported several readings and that more information would not have settled them. The qualification matters, because most claims that analysis would have worked are made about situations that were uncertain in Weick's sense, and for those the chapter's claim does not apply and analysis is the right tool. Where to look: Maitlis and Christianson's review of the sensemaking literature, and the crisis studies of Maitlis and Sonenshein; the falsifier the two reviews sharpen is a case in which a group that committed publicly to an account and kept it open to the cues of its own action did better than one that held its picture apart until it was accurate, where the chapter expects the reverse.²³

A second finding would weaken the chapter without refuting it. The seven properties are a description. If they could be reduced to fewer without loss, or if one of them (the social property, say) turned out not to hold for the solitary experts Klein studied, the description would need revising and the book's use of it would need to be narrower. The chapter has used the properties as a diagnostic, for saying which organizational habit misreads which part of the activity, and that use survives some revision of the list.

What the soldiers kept updating

The unit in the Alps had two things when it was lost: a map that was wrong and a practice of updating that was right. The chapter has been about the second, as an activity: social, retrospective, felt, driven by plausibility, and tested by walking. It has left one question open on purpose. Weick says the soldiers kept updating their sense of where they were. Updating what, held where, and against what? Awareness in chapter 3 and sensemaking here run on the

²³Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21–49. p. 21: “four distinct forms: guided, fragmented, restricted, and minimal. These forms result from the degree to which leaders and stakeholders engage in ‘sensegiving’”; p. 24: three British symphony orchestras; p. 32 (Fig.): guided → “Unitary, rich account”. The two dimensions are control and animation, so “how much leaders and stakeholders control them” was a simplification, corrected in the text.

same thing: the account of the situation and of oneself that chapter 1 called the internal model, and that the map, once corrected, became. How that account is kept current, what it is for, and how far the claim that a regulator needs one can be pushed, is the subject of chapter 5.

References

- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Boje, D. M. (1991). The storytelling organization: A study of story performance in an office-supply firm. *Administrative Science Quarterly*, 36(1), 106–126.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Czarniawska, B. (1997). *Narrating the Organization: Dramas of Institutional Identity*. University of Chicago Press.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64.
- Gephart, R. P., Topal, C., & Zhang, Z. (2010). Future-oriented sensemaking: Temporalities and institutional legitimation. In T. Hernes & S. Maitlis (Eds.), *Process, Sensemaking, and Organizing*. Oxford University Press.
- Gioia, D. A. (1986). Symbols, scripts, and sensemaking: Creating meaning in the organizational experience. In H. P. Sims, Jr. & D. A. Gioia (Eds.), *The Thinking Organization* (pp. 49–74). Jossey-Bass.
- Klein, G., Moon, B., & Hoffman, R. R. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, 21(5), 88–92.
- Maitlis, S., & Christianson, M. (2014). Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals*, 8(1), 57–125.
- Maitlis, S., & Sonenshein, S. (2010). Sensemaking in crisis and change: Inspiration and insights from

- Weick (1988). *Journal of Management Studies*, 47(3), 551–580.
- Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247.
 - Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
 - Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge.
 - Stigliani, I., & Ravasi, D. (2012). Organizing thoughts and connecting brains: Material practices and the transition from individual to group-level prospective sensemaking. *Academy of Management Journal*, 55(5), 1232–1259.
 - Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
 - Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.
 - Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization Science*, 16(4), 409–421.

Chapter 5

Chapter 5 — Orientation: keeping the model alive

A plan that keeps aborting

On page 183 of the second edition of *Brain of the Firm*, in the middle of an argument about the part of a firm that looks ahead, Stafford Beer asks a question that reads at first like a provocation: “a corporate plan must continuously abort, for unless that plan is continuously aborting, how shall sanity (the match between the world and the firm’s notion of it) be preserved?”¹ Chapter 1 quoted the sentence. What it asks for is the subject here.

A plan that aborts has met something it did not contain. What it met was the world. A plan cannot notice anything, so what noticed the mismatch was the firm’s notion of the world, which Beer on the same page calls a model, and of which he says that it “may be disseminated, in separate chunks of cognition around the firm, rather than being cohesive and well formulated”.² Sanity, in his definition, is the match between the notion and the world. Keeping the match is the activity this chapter calls orientation, after John Boyd. The thing kept matched is the internal model, which chapter 1 introduced, which chapter 4 showed breaking and being repaired, and which is developed here.

¹Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

²Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

Orientation at the center of the loop

Boyd's loop, observe–orient–decide–act, is usually remembered as a cycle to be run faster than an opponent's. Frans Osinga, whose study of Boyd is the book's source for him, shows that Boyd put orientation at the center.³ Boyd's essay *Destruction and Creation* describes adaptation as the unstructuring of old concepts and the building of new ones, and in his notes he calls these two processes “the engines for matching observed phenomena or physical patterns with mental models or mental patterns”.⁴ Orientation is that matching, done continuously. It is where the mental models sit, where observation is read, and where the decision comes from. In Osinga's reconstruction of the full loop, decision is a hypothesis and action is its test, and “the entire OODA loop is a double-loop learning process”, because what an action returns is used to check the orientation that produced it.⁵

For organizations the reading changes what the loop is for. Chapter 1 drew Boyd's loop as the inner circle of the book's loop, the one that can turn fast, and the temptation is to build the organization around its speed: shorter reporting lines, faster decision meetings, quicker execution. Boyd's own emphasis points elsewhere. If orientation is the center, the question to ask of an organization is where its orientation is done, by whom, with which models, and how those models get revised. An operating model usually names a place for observing (the reporting), a place for deciding (the management meeting) and machinery for acting. It rarely names a place where the model is revised, and where it has one it rarely calls it that. Speed, as chapter 6 argues, falls out of the quality of the model.

³Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, ch. 3 (*Destruction and Creation*; the “engines for matching” from Boyd's notes as quoted there). The PDF has no page numbers and is cited by chapter.

⁴Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, ch. 3 (*Destruction and Creation*; the “engines for matching” from Boyd's notes as quoted there). The PDF has no page numbers and is cited by chapter.

⁵Osinga (2007), ch. 6, fig. 6.1 “The real OODA loop” and the surrounding text.

Pressing for it directly produces a fast loop on an unrevised picture, which is the organizational form of what the smokejumpers did at Mann Gulch.

The earlier book from which this one grew presented Boyd's loop in the frame of speed, as an answer to a gap in tempo; this book departs from that reading, and Appendix A records the departure.⁶

The internal model, developed here

Chapter 1 took the idea from Kenneth Craik's suggestion of 1943 that thought works with a "small-scale model" of external reality, and used the term for the account an organization acts on, wherever it is held.⁷ Four things about the model need saying before the chapter can use it: what it is for, what it is not, what it must be of, and where it lives.

What it is for. Beer is the bluntest of the book's sources: "the regulator is as good as its contained model of the system regulated, and no more".⁸ Earlier in the same book, in one of the dialogue passages of *The Heart of Enterprise*, he had put the point in the first person: "all we're controlling is the projection of the model-in-our-heads".⁹ Steering is done with the model, and whatever the model does not contain cannot be steered for. That is a claim about necessity, and it goes beyond what Weick, Hutchins or Nonaka say. Weick describes how accounts are built, Hutchins how systems compute, Nonaka how knowledge is created; Beer says that without a model there is no steering at all, and that steering is bounded by the model's quality. How much of that claim can be warranted is the question of a later section. Its use, whatever its warrant, is that it turns

⁶Mathis, K. (2025). *The Tautai Principle*, p. 78 n. 40 (OODA as the primary solution to the speed gap); the departure is recorded in Appendix A of this book and in the edition triggers of the outline.

⁷Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.

⁸Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

⁹Beer (1979), p. 80.

orientation from a nice-to-have into the thing on which everything downstream depends.

What it is not. The model is not the data. Beer italicizes the word to make the point: “*Model* is italicized, because it must not be confused with the variety in the information that happens to be flowing”.¹⁰ His example, which chapter 2 uses, is a pair of parents who cannot recognize the symptoms of their son’s drug use: the information flows past them, and they have no model in which it means anything.¹¹ Ross Ashby had put the general form a generation earlier: “a set’s variety is not an intrinsic property of the set: the observer and his powers of discrimination may have to be specified if the variety is to be well defined.”¹² The model decides what can be seen at all. A cue that fits no distinction the model makes is invisible rather than weak, and the distinctions were drawn by someone. Chapter 2 treats that as the first design decision of perceiving, and chapter 10 as a question of power. Here the consequence is for orientation: revising the model can mean adding a distinction, and a unit whose model has no place for a kind of event will not revise on account of it, however many such events occur.

What it must be of. Beer requires two models where most management vocabulary has one. The part of the organization that looks ahead must contain a model of the corporation; it “must contain a model of the viable system of which it is the System Four”; and the part above it “must necessarily contain models of System Three and Four”.^{13,14,15} A model of the environment on its own is a forecast. Steering needs, in addition, an account of what the steering thing is, what it can do and what it is for. This is where Weick’s first property, identity, gets its structural

¹⁰Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

¹¹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

¹²Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

¹³Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

¹⁴Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

¹⁵Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

place. An organization that has lost its account of itself cannot orient, however good its market intelligence, and chapter 4 showed what that looks like in a crew. Chapter 9 develops the two models as the core of Beer's System Four, following *Punch and Judy* on what System Four is and is not, and does not restate what that book says.¹⁶

Where it lives. The model "may be disseminated, in separate chunks of cognition around the firm"; "no one may think of it, or refer to it, as a model, but it must be there".¹⁷ Raúl Espejo and Alfonso Reyes go a step further and read the structure itself as a model: the organization's structure "should map" the complexity it chooses to respond to.¹⁸ Edwin Hutchins' navigation team held its model of the ship's position in a chart, a bearing log and a set of routines as much as in anyone's head, as chapter 1 showed. The chunks, then, are of three kinds: heads, artifacts and structure. That is the reason the model can be designed, which is Part III's claim, and the reason it can be lost without anyone noticing, which is Part IV's.

Planning forward, sensemaking backward

Gary Klein watched experienced people do something he called mental simulation, "the ability to imagine people and objects consciously and to transform those people and objects through several transitions", and he found them doing it in two directions: forward, to test a course of action before committing to it, and backward, to explain an event that made no sense.¹⁹ Mica Endsley and Debra Jones make mental models "the key enablers of Levels 2 and 3

¹⁶*Punch and Judy*, ch. 2 "System Four Is Not a Pair of Eyes", section "Action selection: the function the diagram omits" (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

¹⁷Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

¹⁸Espejo, R., & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Springer, p. 95; see also p. 138.

¹⁹Klein (1998), ch. 5 "The Power of Mental Simulation" (from p. 49).

SA”, that is, of comprehension and projection, the two levels of awareness that go beyond perceiving.²⁰ Beer, on the page this chapter opened with, calls foresight “rapid-time simulation”.²¹ And Andy Clark, writing from cognitive science, describes brains as “prediction machines” that run a “generative model” of what they expect to perceive.²²

Four sources, one claim: there is one model, and it is used in two directions. Run forward it is planning, projection and, in chapter 4’s terms, prospective sensemaking. Run backward it is explanation and sensemaking in Weick’s sense, the repair of a frame that did not predict what happened. Chapter 6 adds a third use, forward on a candidate act, which is how a recognized action is tested before it is released. The three uses draw on the same account of the world and of oneself, and a fault in the account shows up in all three. That is the practical reason for treating them together: an organization that plans on one model and explains its failures on another has two models, and neither is being corrected by the other.

The order of the two directions matters for a design. Most organizations run the model forward on a calendar (the budget, the plan, the roadmap) and backward on demand, after something has gone wrong. Boyd’s loop and Klein’s experts do both continuously, and the backward run is what keeps the forward run honest. Where the backward run is missing, the forward run degrades into extrapolation.

Clark’s vocabulary is tempting here, and *Punch and Judy* sets the limit on its use. The predictive-processing picture of a brain that continuously predicts its input and revises on the error fits Beer’s System Four closely. It is a comparison, not Beer’s claim. No passage in Beer’s three main books calls the line from System Four to command a prediction channel, and the term forward model does

²⁰Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press, pp. 21–22.

²¹Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

²²Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, pp. 3–4.

not occur in them.²³ Predictive processing is a theory of brains; carrying it over to organizations is an analogy, and the book marks it as one whenever it appears. The free-energy principle that underlies the more ambitious versions is examined, with its critics, in *Punch and Judy*, and the book does not lean on it.²⁴

Navigating with a model, not following a plan

Beer's sentence about the aborting plan draws the distinction on which the chapter's title rests. A plan is a model run forward, frozen at a date, with the assumptions that produced it usually left out. A model is what produced the plan and what can notice that the plan no longer fits. Having a model and following a plan are therefore different activities, and an organization can do the second without the first. It is the second that Beer calls insanity, in his definition of sanity as the match between the world and the firm's notion of it.

Two objections from outside cybernetics say the same thing from the other side. Lucy Suchman argued that plans are resources for action rather than programs that control it; Rodney Brooks, that "the world as its own model".^{25,26} Both were writing against the idea that intelligent behavior is the execution of a stored representation, and both warn against mistaking a model for a plan. That is Beer's warning too. The disagreement between them and Beer is about how much model an actor needs to carry, and the book comes to that in the section on the two senses.

²³*Punch and Judy*, ch. 2 "System Four Is Not a Pair of Eyes", section "Action selection: the function the diagram omits" (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

²⁴*Punch and Judy*, ch. 12, on the free-energy principle and its critics.

²⁵Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.

²⁶Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159. The paper says "the world as its own model" (p. 139), not "the world is its own best model".

The earlier book called the habit of following a plan as if it were the model a prediction addiction, and set navigating against it.²⁷ The navigator of chapter 1 is the image to keep. The navigators of the Caroline Islands had a model, the *etak*, a way of organizing what the navigator knows, updated segment by segment as the voyage goes on. Navigating is acting with a model that keeps being corrected. Planning, in the sense the chapter is against, is acting on the model's last output as if the model had stopped.

What has to be true for a plan to abort when it should? Someone has to notice the mismatch, which needs the plan's assumptions to be stated somewhere, since a mismatch is a mismatch with something. That person has to have a forum in which the mismatch can be said, and saying it must not count as disloyalty. There has to be a record of what was expected, which is chapter 6's decision record. And the abort has to lead somewhere: to a revised model and a new forward run, rather than to a search for who was wrong. None of these is difficult to install. All of them are routinely absent, because a planning cycle is built to produce the plan and to defend it, and the abort is an admission that the cycle was not built to make.

The two senses of the model, and the warrant

The book uses the word *model* in two senses, and this is where it states them and says which it uses where.

In the *loose* sense, every organization acts on some account of its situation and of itself. That account is held in Beer's separate chunks of cognition, in heads, artifacts and structure, and the only questions are where it lives, who drew its distinctions, whether anything decided depends on it, and how it is revised. The sense is uncontroversial. Nobody denies that a firm has a notion of its market, and nobody claims that the notion is written down in one place. When the book says, in chapters 1 to 4 and again in Part III and Part IV, that a unit's model filtered what it saw, or

²⁷*The Tautai Principle*, pp. 77–78, as listed in the outline for this chapter. Appendix A gives the mapping of navigating and planning to the earlier book's terms.

broke, or was repaired, or was lost with the person who held it, it means the loose sense, and no more is claimed.

In the *strong* sense, a regulator needs a model of what it regulates, and it steers no better than that model. This is Beer's claim in the sentence quoted above, and he called it a "cybernetically proven fact".²⁸ The proof he had in mind is Roger Conant and Ross Ashby's paper of 1970, whose title is "Every good regulator of a system must be a model of that system". The cybernetic tradition has cited that title as a theorem ever since. *Punch and Judy*, which this book follows on every point where the two touch, went back to the paper, and the book reports what it found rather than repeating the citation.²⁹

The paper does not prove its title. Conant and Ashby define a good regulator as one that minimizes the entropy of the outcomes, which means that a regulator pinning the system to one constant catastrophic outcome counts as optimal; their good regulator is "a *reliable* regulator, not a desirable one". The authors concede in the paper that not all optimal regulators are models. The proof mentions no inputs, no channels and no noise, and it holds only where the disturbances change "slowly with time", which excludes exactly the environments a System Four exists for.³⁰ The rigorous form of the claim is Bruce Francis and Murray Wonham's internal model principle of 1976, which

²⁸Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

²⁹*Punch and Judy*, ch. 3 "The Law We Quote Is Not the Law Ashby Wrote", section "The theorem that is not a theorem", on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors' concessions at p. 96. Beer's "cybernetically proven fact" at *The Heart of Enterprise*, p. 114.

³⁰*Punch and Judy*, ch. 3 "The Law We Quote Is Not the Law Ashby Wrote", section "The theorem that is not a theorem", on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors' concessions at p. 96. Beer's "cybernetically proven fact" at *The Heart of Enterprise*, p. 114.

holds for control problems with stated dynamics; an organization is not obviously such a problem.³¹ A second route is the law of requisite knowledge: having enough possible actions is not enough, since the regulator must also know which action fits which disturbance, and that knowledge is a model of sorts.³² A third is the reconstruction by Nathaniel Virgo and colleagues, which replaces *being* a model with *having* one: an observer can attribute to a regulator beliefs that update consistently with the system it regulates, at the price that the attributed model may be trivial.³³

The price is worth paying, because it makes the strong sense checkable. *Punch and Judy* turns the familiar complaint that a planning unit scans and nothing downstream moves into a claim that can be tested: the model attributable to this regulator is trivial, since nothing the organization decides depends on it.³⁴ That is the diagnostic question the book carries from here: which model does this unit steer by, of what, and does anything it decides actually depend on it?

So the strong sense is argued for, within limits, and not proved for organizations. The book uses it in four places, and in each it says so. Here, in the claim that orientation quality bounds steering, which is what makes orientation the center of the loop. In chapter 6, in the claim that the speed of a recognized action is a consequence of the model's quality. In chapter 9, in the diagnosis of a System Four whose attributed model is trivial. And in chapter 13, in the claim that delegation to a machine can go only as deep as the capacity to evaluate what comes back, since evaluating presupposes a model of the task. Everywhere

³¹*Punch and Judy*, ch. 3, on Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica*, 12(5), 457–465.

³²*Punch and Judy*, ch. 3, note 14, on Heylighen, F., & Joslyn, C. (2001). Cybernetics and second-order cybernetics. In R. A. Meyers (Ed.), *Encyclopedia of Physical Science & Technology* (3rd ed., vol. 4). Academic Press.

³³*Punch and Judy*, ch. 3, on Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326; also Baez (2016) and Wentworth (2021).

³⁴*Punch and Judy*, ch. 3, section “From a complaint into a claim”.

else the loose sense does the work, and the design grids of chapter 14 rest on the loose sense plus the diagnostic question.

Two further limits come with the warrant. The good-regulator reading, even where it holds, says nothing about whether the state being defended is worth defending; a regulator can be excellent at holding a firm to a purpose that has stopped making sense. That question falls to identity and purpose, and chapter 10 takes it up.³⁵ And the neighboring law, Ashby's law of requisite variety, is used in this book as a design argument only. Ashby wrote that "only variety in R can force down the variety due to D; variety can destroy variety"; the sentence usually quoted, that only variety can absorb variety, is Beer's compression and travels without the law's conditions.³⁶ The law is a priori, "it has nothing to do with the properties of matter", so it licenses an argument about whether an arrangement could work, never a figure.³⁷ The book reports no variety figures and quotes Ashby's own sentence.

Weick's map and Beer's regulator

Chapter 4 ended with the Hungarian unit's map, and the two senses let the book say where it stands between that map and Beer's regulator.

Weick asks of a model only that it get people moving and noticing cues. The soldiers "kept moving, they kept noticing cues, and they kept updating their sense of where they were", and the map's relation to the mountains was beside the point.³⁸ Beer asks that the model fit what it

³⁵*Punch and Judy*, ch. 3 "The Law We Quote Is Not the Law Ashby Wrote", section "The theorem that is not a theorem", on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors' concessions at p. 96. Beer's "cybernetically proven fact" at *The Heart of Enterprise*, p. 114.

³⁶Ashby (1956), p. 207 (section 11/7), pagination as in *Punch and Judy*, ch. 3. "Only variety can absorb variety" is Beer's compression (*Brain of the Firm*, 2nd ed., p. 308).

³⁷*Punch and Judy*, ch. 3, note 5: Ashby (1956), p. 208 (section 11/10).

³⁸Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 54–55.

steers, because steering is no better than the fit. Suchman and Brooks stand nearer to Weick: carry as little model as you can, and let the world supply the rest. The three positions look incompatible, and the book's position is that they describe different moments of one loop.

An orienting model must be plausible enough to act on. That is Weick's requirement, and it is a requirement on the input to the loop: without a first account, however wrong, nothing moves and no cues are produced. A regulating model must fit what it regulates. That is Beer's requirement, and it is a requirement on the output of the loop: what the account has to have become, through action and cues, if steering by it is to be better than steering by nothing. The two requirements are made of the same model at different times, and the loop is what takes it from the one state to the other. Beer's continuously aborting plan is the mechanism: the plan is the orienting model's forward run, the abort is the moment the cues refuse it, and the revision is what moves the model toward fit.

Nonaka's standard for knowledge marks the same distinction from the side of Part IV. He defines knowledge as "justified true belief", and stresses justification as a process rather than a state.³⁹ Plausibility is what is needed to act; justification, in Nonaka and Takeuchi's sense, is the organization's later judgment of whether the account is worth committing to, taken against the purposes it has set. Weick's sensemaking supplies the first and Nonaka's knowledge creation the second, and an organization needs both, in that order.

Beer gives the loop's end state a name that belongs in Part IV as well. When a pupil and a teacher each build a model of the other and "the two models of each other converge, learning has occurred".⁴⁰ Orientation is that

³⁹Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14-37, p. 15. <https://doi.org/10.1287/orsc.5.1.14> (public PDF, pp. 14-21). Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.

⁴⁰Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

convergence, kept up: the organization's model of its environment and the environment's actual behavior brought toward each other by rounds of action and correction, without ever arriving. When the rounds stop, the plan stops aborting, and what is left is Beer's insanity, which from inside feels like a plan being followed well.

The representation chosen

Orientation depends on the representation the model is held in, and chapter 1's navigator is again the evidence. For the Caroline Island navigator, Hutchins reports, "a canoe on course between islands is stationary and the islands move by the canoe"; the voyage is tracked as a sequence of *etak* segments, each marked by a reference island passing under a star bearing.⁴¹ The Western navigator holds the same voyage as a moving point on a fixed chart. The computational task is the same in both traditions. What differs is which updates are easy and which are impossible: the *etak* makes the passage of time and the next bearing the natural units of revision, and the chart makes position and distance the natural units. Neither is more intelligent. Each makes a different set of mismatches visible.

Organizations choose representations all the time and rarely notice that they are choosing. A dashboard, a strategy map, a risk register and a customer journey are each a way of holding the model, and each decides which kind of cue can register as a mismatch and which cannot. A dashboard built around last year's segments cannot show that the segments have moved; a risk register organized by named risks cannot hold a risk that has no name yet. Chapter 1's measurement trap is a case of this. Chapter 7 develops the general point with Hutchins: the chain of representations through which an account travels is organization design, and each link decides what the next one can see.

⁴¹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 2, p. 71, following Gladwin, T. (1970), *East Is a Big Bird*. Page numbers are the printed pages of the calibre scan.

For orientation the consequence is that revising the model sometimes means changing the representation, and that this is the revision organizations resist most, because the representation is built into the tools. It is easier to update a number on a dashboard than to admit that the dashboard's rows are the wrong rows.

Rhythms of re-orientation

How often an organization returns to its model is a design variable, and the chapter states it as one without prescribing a value. Boyd's loop turns continuously, and Klein's experts simulate as a matter of course; that is orientation at the rhythm of the work. Chapter 1's outer circle turns over months and years, and the strategy review, the budget round and the retrospective are its usual forums. Between the two lie the rhythms an operating model sets, mostly without deciding to: how often a unit revisits its account of the market, of the customer, of itself.

Two things can be said about the rhythm without numbers. It has to be shorter than the time it takes the terrain to change, or the model will be corrected after the plan it produced has already failed. And re-orientation on a calendar is only half of it; the other half is re-orientation on interruption, chapter 4's trigger, which no calendar can schedule and which a forum has to be able to receive when it comes. An organization that re-orientes only on the calendar will be surprised between reviews; one that re-orientes only on interruption will drift until something breaks. Chapter 14 makes the rhythm per level of the organization one of its six design grids, and the rhythm differs by level: a team facing its own customers and a company facing its regulators need different ones.

What the rhythm does not settle is speed. A short rhythm of re-orientation is a way of keeping the model good. It is not the same as pressing the loop to turn faster, and the difference between the two is the argument of chapter 6.

What would count against it

The chapter claims that orientation, understood as the continuous revision of an internal model, is the center of the loop, and that the model concept is needed to say what goes wrong when orientation fails.

Two findings would count against it. The first: organizations that revise their models rarely, but hold them accurately, adapting as well as organizations that revise continuously, in comparable environments. If that held, the rhythm of re-orientation would be a matter of taste, and the aborting plan would be one style among others. The chapter's answer is that an accurately held model in a changing environment is a model that is being revised, whether the organization calls it that or not; but the answer is an argument and not a finding, and it should be tested where it can be.

The second: the model concept adding nothing to the diagnosis of orientation failures beyond what situation awareness and sensemaking already give. If every failure of orientation can be placed at one of Endsley's three levels or described as a frame that broke, without ever needing the question which model a unit steers by, then the concept is a name and not a tool. Where to look: Endsley and Jones list "errant mental models" among the eight demons of situation awareness, and their cases are the place to ask whether the demon can be diagnosed without asking what the model was and where it lived.⁴² And *Punch and Judy's* chapters 2 and 3, from which the chapter has taken its warrant, are where a reader should go to see how much of the strong sense survives a closer reading of Conant and Ashby than the tradition has usually given it.

A model that proposes actions

The chapter opened on a plan that must keep aborting and has ended with a rhythm at which it should. What it has not said is what happens when the plan aborts and

⁴²Endsley & Jones (2012), ch. 3, p. 31: the eight demons of situation awareness, "errant mental models" among them.

the model is revised: something has to propose the next action, and something has to release it. Klein's experts recognized the action in the situation, and the same model that made sense of the situation delivered what to do. That is the model run forward on a candidate act, the third use named above, and it is where orientation becomes choosing. Chapter 6 takes it from there, with Klein's commanders, with Boyd's decision as hypothesis, and with the question the whole loop has been circling since chapter 1: whether the point of all this is to be fast.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
- Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press.
- Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press.
- Espejo, R., & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Springer.
- Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica*, 12(5), 457–465.
- Gladwin, T. (1970). *East Is a Big Bird*.
- Heylighen, F., & Joslyn, C. (2001). Cybernetics and second-order cybernetics. In R. A. Meyers (Ed.), *Encyclopedia of Physical Science & Technology* (3rd ed., vol. 4). Academic Press.

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Mathis, K. (2025). *The Tautai Principle*.
- Mathis, K. *Punch and Judy* (author's manuscript).
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14-37.
- Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge.
- Suchman, L. A. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.
- Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.

Chapter 6

Chapter 6 — Choosing: from understanding to commitment

A good course of action from the start

The fireground commanders Gary Klein studied for the work that became *Sources of Power* made “around 80 percent of their decisions in less than one minute”.¹ The commanders arrived at a fire, looked, and acted. When Klein examined how the action had been chosen, he found that in most cases it had not been chosen from among alternatives at all. The commanders “did not have to compare options”; they “could come up with a good course of action from the start”.²

What they were doing instead was recognizing. A commander saw a fire of a kind he had seen before, and with the recognition came an action that had worked for that kind. Klein’s name for the other procedure, in which options are laid out and weighed against each other, is “comparative evaluation”.³ In the fireground data it was rare, and where it appeared it was mostly novices who used it; the clearest case in the study came from a novice at a fire in an oil tank farm.⁴ In a later study of naval officers the share of recognitional decisions was 95 percent.⁵ Klein

¹Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press, p. 4.

²Klein (1998), ch. 3, about p. 17 (page approximate; the offset was checked for pp. 4, 12, 20 and 24).

³Klein (1998), p. 12 (definition of comparative evaluation).

⁴Klein (1998), p. 24 (“fully 80 percent” recognitional; the novice case at the oil tank farm; comparison for novices, for justification and for novel situations).

⁵Klein (1998), p. 103 (the naval study, 95 percent recognitional).

connects this to Herbert Simon's satisficing, "selecting the first option that works": the commander does not look for the best action, because looking costs time and the first workable action is usually good enough.⁶

The recognition is not the end of the process, only its visible part. Where the commander was unsure whether the recognized action would work, he ran it forward in his head and watched what happened: Klein calls this mental simulation, "the ability to imagine people and objects consciously and to transform those people and objects through several transitions", and devotes a chapter to it.⁷ Chapter 5 described the internal model as something run forward in planning and backward in explanation. Mental simulation is the model run forward on a candidate action, and it can reject the candidate before anyone has moved. Klein found the same faculty used in the other direction, to explain a strange event by imagining how it could have come about, which is the backward use chapter 5 described.⁸ The test is cheap, it is fast, and it is done by the same person who did the recognizing, with the same model. That is the arrangement most organizations dismantle when they formalize a choice: the person who recognized the situation writes a proposal, and a different person, with a thinner model of the situation, is asked to evaluate it against alternatives that the first person never seriously entertained.

For this book the finding does two things. It closes the gap that chapter 1 left open when it said that understanding comes first. The five activities of that chapter, drawn as a loop, invite the reading that choosing is a stage that begins when understanding is done. For the commanders it was no such stage. The model that made sense of the fire proposed the action and tested it; the choice was part of the understanding, and the separation of the two is something novices and novel situations need. Chapter 1 stated this as a qualification. Here it becomes the starting point: planning, sensemaking and choosing are three uses of one

⁶Klein (1998), p. 20, on Simon's satisficing as Klein cites it.

⁷Klein (1998), ch. 5, "The Power of Mental Simulation", from p. 49.

⁸Klein (1998), ch. 5, "The Power of Mental Simulation", from p. 49.

model, forward, backward, and forward on a candidate act.

The second thing the finding does is set a limit, which the rest of the chapter has to respect. Klein did not find that experienced people never compare. Comparison is what novices do, what justifying a choice to other people requires, and what a situation that fits no known type demands.⁹ And recognition has a failure mode of its own, which the commanders' successes do not show: a situation that only looks familiar delivers an action with the same confidence as one that is. The commander who recognizes the wrong type will act fast and wrong. That case returns later in the chapter, when Cynefin is placed.

Every action is also a question

John Boyd is usually cited for a loop of four verbs, observe, orient, decide, act, and for the claim that the side which cycles faster wins. Frans Osinga's reconstruction of the full loop gives a different picture.¹⁰ In the full loop, decision corresponds to hypothesis and action to test. Decisions "feed back into the systems as validity checks on the correctness and adequacy of the existing orientation patterns", and, in Osinga's summary, "the entire OODA loop is a double-loop learning process".¹¹ The loop is a way of finding out whether the orientation is right, and speed is one of its properties, which the chapter comes back to. (Chapter 11 sorts out the several meanings that "double loop" has acquired; Osinga's is one of them.)

Andy Clark arrives at the same place from the cognitive science of prediction. Perceiving, in his account, is

⁹Klein (1998), p. 24 ("fully 80 percent" recognitional; the novice case at the oil tank farm; comparison for novices, for justification and for novel situations).

¹⁰Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, ch. 6, fig. 6.1 "The real OODA loop" and the surrounding text. (the PDF has no page numbers).

¹¹Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, ch. 6, fig. 6.1 "The real OODA loop" and the surrounding text. (the PDF has no page numbers).

“intimately bound up with various forms of motor planning and action selection”; an agent that holds a hypothesis about the world can use “cheap epistemic action” to “confirm or disconfirm the hypothesis”.¹² He separates two kinds of action: pragmatic actions, which move the world toward a goal, and epistemic actions, which change what the agent knows.¹³ The distinction comes from David Kirsh and Paul Maglio’s study of Tetris players, which is the standard reference for it.¹⁴ Karl Weick, from the other side of the book’s theories, says it in three words: “Action shapes cognition.”¹⁵

Every action, then, has two values. It achieves something, and it reveals something, and the two do not have to be in proportion. A pilot test in one region, a prototype shown to five customers, a price changed for a month: each does a little to the world and tells the organization a great deal about whether its picture of the world holds. The belief core quoted in chapter 1 asks for “probes with stop conditions rather than plans with milestones”.¹⁶ Clark’s distinction gives the probe its definition: a probe is an action chosen mainly for its epistemic value, and the stop condition is the point at which the hypothesis is declared confirmed or refuted. Choosing is therefore part of understanding twice over. Recognition delivers the action out of the model, and the action delivers the next revision of the model.

Mann Gulch, read this way, shows what happens when the two come apart. The crew’s recognition of the fire as a familiar type delivered the familiar action, and the type was wrong. Wagner Dodge’s escape fire was an action of

¹²Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, p. 183 (“cheap epistemic action”; “confirm or disconfirm the hypothesis”).

¹³Clark (2016), p. 184 (pragmatic and epistemic action).

¹⁴Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549. pp. 514–515 (pragmatic vs epistemic action), p. 527 (Tetris rotations save mental rotation effort).

¹⁵Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, p. 12.

¹⁶The Tautai company belief core, 01 Tautai/01-Canon/01. Introduction/00.core-belief/00.core-belief.md (technology and authority). Standalone readers: Appendix A.

the other kind: invented, not recognized, and its meaning could not be constructed by the others in the time they had.¹⁷ An action whose epistemic value nobody else can read is a private experiment, and it may save one person.

Most of what could be done is not done

Between the model and the act there is a step that the loop diagrams leave out, and *Punch and Judy*, a book on the Viable System Model that this book follows on cybernetic questions, calls it “the function the diagram omits”.¹⁸ At any moment an organism, or an organization, has more actions available than it can perform, and most of them are incompatible with each other: one cannot advance and retreat, hire and freeze, launch and delay. In the words of that chapter, “something has to decide which one is expressed”. In the vertebrate brain, as *Punch and Judy* reports it, that something is a set of structures whose output is inhibitory by default: every candidate action is held back, and selection happens by releasing one from inhibition. The central problem becomes “how a field of possible actions is held back, compared, and selectively released”.¹⁹ The anatomy is borrowed; I use it as an analogy and nothing more.

¹⁷Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633 (cosmology episode: people “suddenly and deeply feel that the universe is no longer a rational, orderly system”; minimal organizations “susceptible to sudden losses of meaning”), p. 635 (drop-your-tools as a loss of identity). The reading of the escape fire as an invented action whose meaning could not be constructed in time is mine; the case is told in chapter 1.

¹⁸*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

¹⁹*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

The analogy is useful because it reverses the usual picture. Organizations describe their problem as generating options: more ideas, more initiatives, more experiments. The scarce capability is more often the opposite. It is suppressing most of what could be done and releasing little of it, and releasing it decisively. Several familiar management devices are release-and-inhibition mechanisms under another name. A stop condition attached to a probe is an inhibition set in advance: the action is released for a defined time and returns to the held-back field unless the condition is met. A work-in-progress limit inhibits the release of new work until old work is finished. Portfolio pruning is inhibition applied late, to actions that were released and never stopped. A review rule from the consulting practice this book grew out of allows an experiment to “continue watching” at most once per cycle, on the ground that “indefinite watching is how zombie experiments accumulate”.²⁰ The rule exists because the default in most organizations is the reverse of the vertebrate one: once released, an action stays released until someone finds the energy to stop it.

The result is visible in any portfolio review. The list of running initiatives is long, each has an owner who can explain why it should continue, and the meeting spends its time on which new items to add. Nobody is charged with inhibition, so nothing is inhibited, and the organization’s capacity is spread across a field of actions that were each released for a reason and are now each starved of the attention that would let any of them finish. Release without inhibition produces the same outcome as no selection at all, only slower and at greater cost. The inhibitory default of the analogy is worth taking literally as a design principle: an action should have to earn its release, and should lose it again by default.

Attention, in chapter 2, was the selection problem on the input side: of everything that could be noticed, what gets through. Action selection is the same problem on the

²⁰Tautai cycle review, 01 Tautai/11-Consultant Support/20-Event-Kits/phase-4-intervene-experiment/cycle-review.md.

output side. Both are selection under scarcity, and both are mostly done without anyone deciding. Chapter 9 notes that the Viable System Model's diagram has no place for the second of them, and asks where it should go.

Two words for two things

Most selection in an organization is continuous and tacit. A service agent decides how to answer a complaint; an engineer decides which of three bugs to fix first; a team lead decides to let a late deliverable slide another day. None of these is minuted, and each is a choice, made out of a model of the situation, of the kind Klein's commanders made. The book calls this *action selection*. A *decision*, in the sense the book reserves for the word, is different in kind: an explicit, accountable commitment of resources that other people can rely on and later check.

The reason for keeping the two words apart is not tidiness. It is what happens to a commitment that has no record. Weick, drawing on Harold Garfinkel's study of jurors, observes that people often decide first and construct the grounds afterwards: "the person defines retrospectively the decisions that have been made".²¹ Garfinkel's jurors, in Weick's reading, arrived at a verdict and then assembled the reasoning that justified it. Organizations do the same with their commitments. Six months after a reorganization nobody can say what it was supposed to achieve, because the expectation was never written down, and each participant now remembers a different one, shaped by what has since happened. The commitment has been redefined in retrospect, and there is nothing left for learning to test against.

A record changes this. Sensemaking Article 2, written for this book's practice, proposes a "commitment forum" whose output is a "decision record, owners, metrics, review date".²² Article 3 treats such records as organizational

²¹Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 10–11, on Garfinkel's jurors.

²²Sensemaking Article 2, 51-Stories/Sensemaking Articles/article-02-sensemaking-moves.md, "Blueprint C —

memory: the place where the organization keeps what it once expected, so that it can later find out whether the expectation held.²³ Read with Osinga, a decision record is Boyd's hypothesis in written form. It states what the organization believes about its situation, what it is committing to on that belief, and what would show the belief to have been wrong. The next round of sensemaking then has something to confirm or refute, and the learning of Part IV has material to work with.

The forum matters as much as the record, and for a reason that chapter 4 prepared. When the picture has broken, several accounts of the situation are in play, and the work of comparing them is sensemaking. A commitment is a different act: it ends the comparison, for now, and binds resources to one account. Organizations that do both in the same meeting do both badly. The comparison is cut short by whoever has the authority to commit, and the commitment is weakened by the accounts that were never fully heard and now continue underground. The commitment forum of Article 2 separates the two: the competition of accounts happens first, in a setting where no commitment is yet on the table, and the commitment is made afterwards, on the account that survived, with the record naming what it rests on.²⁴ Chapter 8 returns to the forum as a collective mechanism and to the conditions under which the competition of accounts is honest. Here the design point is the separation itself. I would add that the record should carry the losing accounts as well: what the organization decided against is part of what it can later learn from.

Commitment forum". Article 3 on decision records as organizational memory.

²³Sensemaking Article 3, 51-Stories/Sensemaking Articles/article-03-representational-governance.md, "Decision memory (decision records)".

²⁴Sensemaking Article 2, 51-Stories/Sensemaking Articles/article-02-sensemaking-moves.md, "Blueprint C — Commitment forum". Article 3 on decision records as organizational memory.

The classic organizational account of decisions of this second kind is Henry Mintzberg, Duru Raisinghani and André Théorêt's study of unstructured decision processes.²⁵ The naturalistic decision-making tradition that grew out of Klein's work has since collected the first kind, in the field, across many domains.²⁶ The two words are meant to let this book use both.

The failure modes differ, and chapter 1 named them. Action selection fails when the model behind the recognition is wrong; the remedy is a better model, which is a matter of experience, feedback and the learning of Part IV. Decision fails when commitments are made without a record; the remedy is a forum and a document, which is a matter of design. An organization that treats every selection as a decision, and demands option comparison and sign-off for each, is slow for no gain. One that treats every decision as a selection, and lets commitments happen in corridors, cannot learn from them.

Who selects

Stafford Beer, in *Brain of the Firm*, borrows a term from the neurophysiologist Warren McCulloch for what happens in a nervous system when no fixed center is in charge: the "redundancy of potential command". In such a system "the command centre changes from time to time. Its location is a function of the information available".²⁷ Beer applies it to the firm without apology: "Who, for example, really decides to buy ... It is the little group of employees, possibly very junior, which understands the need". And

²⁵Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The structure of "unstructured" decision processes. *Administrative Science Quarterly*, 21(2), 246–275. p. 246: a field study of 25 strategic decision processes; "Unstructured refers to decision processes that have not been encountered in quite the same form and for which no predetermined and explicit set of ordered responses exists in the organization"; p. 251: 14 of 25 decisions had "custom-made" solutions.

²⁶Zsombok, C. E., & Klein, G. (Eds.) (1997). *Naturalistic Decision Making*. Erlbaum.

²⁷Beer (1981), p. 232 (McCulloch's redundancy of potential command).

he draws the structural consequence: in a rigid pyramid “there can be no redundancy of potential command”.²⁸

Beer’s argument is the one this chapter has been building toward from the other side. If understanding is distributed, as Part III will argue, and if for experienced people understanding already carries the action, as Klein found, then the action is already being selected where the relevant model sits, whatever the organization chart says. The little group that understands the need has in effect decided to buy; the formal decision, weeks later, ratifies a selection that has already shaped every step since. Decision rights placed elsewhere do not move the selection. They add latency, and they add the retrospective redefinition that Weick describes, because the ratifying body has to construct grounds for a choice it did not make.

Hutchins offers a case at the far end of the scale, where the choice was made by no one. When the *Palau* lost its gyrocompass, the navigation team’s division of labor changed, and the new arrangement was, in his analysis, “evolutionary in the sense that the system-level change that resulted was never represented”.²⁹ Nobody proposed it and nobody approved it; it came about in the adjustments the members made to each other’s work under pressure, and it held the ship. That is action selection at the level of the system, without a decider. It is also, as chapter 1 said, the reason the solution was lost afterwards: what was never represented could not be kept.

The argument has a structural version, which chapter 9 develops and which is stated here because it belongs to choice. Chapter 1 said that every unit facing its own environment has to produce its own understanding. It follows that every such unit selects its own actions, whether or not the organization chart grants it the right, because the actions come out of the model and the model is there. What a higher level can add is a commitment point with a record, and a check on whether the lower level’s model is one it should be acting on. What it cannot do, from a

²⁸Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

²⁹Hutchins (1995), ch. 8, p. 349.

distance and with a thinner model, is select better. Beer's little group of junior employees is a unit of this kind, and, extending his remark, a structure which denies it the right to command has not removed the command; it has removed the possibility of the command being recognized and recorded where it happens.

Two design conclusions follow, and they pull in different directions. The first is that decision rights should sit where the model sits: the people who can recognize the situation are the people who can select the action, and moving the right away from them buys accountability at the price of quality and time. The second is that selection at the system level, however good, leaves no trace, and an organization that wants to learn from its choices needs a point at which some of them become decisions in the book's sense, with a record. Where to put that point, and how many selections to pass through it, is a design variable, and chapter 14 treats it as one.

Episode — Eleven weeks, then six days

In the last week of February the work log of a C4 controller in region South carried a line that technicians in all six regions of Halvern's After-Sales Service would go on using until May: "parameter check, 25 min." What Tobias, the field technician, had done in those 25 minutes was raise the minimum compressor off-time and switch off the pass-through of setpoint overrides in the gateway to the building-management system. The two settings came from Milan, a senior second-line specialist, who had worked them out on escalated cases earlier that month and told Tobias on the phone what to change.

Milan had no right to issue them. Only the Product Quality Board, which met monthly under the COO, could approve an official service bulletin, and in February the Board's account of the winter's compressor failures was still installation errors at partners. So

the settings traveled through the technicians' messenger group, first in the South and then in the other regions. A technician standing in front of a short-cycling unit with a building-management system behind it applied them. Nobody above the second line had decided anything, and the action had been selected all the same, by the people who could recognize the fault. The log entry was where the distance showed. "Parameter check" was a term the form already used for routine visits, so it drew no questions, and it left no way to tell afterwards which units had been changed, on whose advice, or what the technician had expected to happen. In March Milan was on leave for three weeks. Technicians whose units did not respond to the two settings had nobody to ask, and their cases stalled. In the same month Katrin, the Head of Service, took Milan, Tobias and three traced cases to the Board, and the account changed to an interaction fault between C4 firmware and the overrides.

The official bulletin was approved in May, eleven weeks after Milan had the workaround. Its rationale rested on the three traced cases and on tests that Quality ran in April, and it presented the two settings as a measure adopted by the Board. The units changed since February did not appear in it, because nothing in the records identified them. When the Board asked in the summer whether the measure had worked, the failure figures before and after May could not answer: an unknown share of the "before" units had been running on the new settings for weeks.

After the winter the rights were moved. The second line could now issue a provisional service bulletin, valid for 30 days, with a record naming the author, the conditions under which it applied and the effect expected; the Board reviewed it afterwards and confirmed, amended or withdrew it. That autumn a faulty sensor batch went from the first report to a provisional bulletin in 6 days. When the Board met, it had a dated

document in front of it instead of a selection it had to reconstruct, and one question left to ask that nobody had been in a position to put to the February workaround: whether the second line had recognized the right kind of fault.

Where recognition can be trusted

Chapter 1 introduced Dave Snowden's Cynefin framework for a first purpose: to say which situations need sensemaking at all. It returns here for a second, because the domain of a situation also decides what recognition can be trusted to do.³⁰

In Snowden's complicated domain, cause and effect are knowable but not obvious, and the right response is analysis by people who know the field. In the complex domain, patterns emerge from the interaction of many agents; they can be understood in retrospect and not predicted, and the right response is to probe, sense and respond. Klein's commanders worked in situations that were dangerous and fast, and that also belonged to types a career of fires had taught them to see. Their recognition had been formed against feedback, fire after fire. Snowden and Boone's warning is aimed at situations of the other kind, where what looks like a type is a pattern that has emerged once and may not recur, because in their complex domain patterns can be read only in retrospect. That recognition there delivers an action with full confidence and no warrant is this book's own reading of their argument, not their statement.³¹

³⁰Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76. Klein's and Snowden's positions are as summarized in the theories table of chapter 1; that they "disagree about how often each case occurs" is my reading, not a documented exchange.

³¹Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76. Klein's and Snowden's positions are as summarized in the theories table of chapter 1; that they "disagree about how often each case occurs" is my reading, not a documented exchange.

The two positions are usually presented as a quarrel, and the book keeps both. Klein's evidence is about experienced people in domains where experience can be formed; Snowden's argument is about domains where it cannot, or not yet. The practical question is which domain a situation is in, and that question is itself a matter of understanding. In my reading, a good part of the damage done by confident recognition in organizations comes from treating a complex situation as a complicated one: recognizing a market, a merger or a technology as a type when it is an emergent pattern. The remedy is the one already described. Where recognition cannot be trusted, an action chosen for its epistemic value, with a stop condition, is the appropriate mode, and the decision record states the hypothesis the probe is testing. Where it can be trusted, demanding probes is waste. On this reading, Klein and Snowden disagree about how often each case occurs rather than about what to do in each.

Speed

The commanders were fast. Around 80 percent of their decisions in under a minute, and the book has now said why: the model was good, and a good model delivers the action with the recognition. Nobody on the fireground pressed for speed. Speed was what a well-formed model looked like from the outside.

That order matters for organizations, because most of the management writing on loops of this kind runs it the other way. It takes the speed of the loop as the goal and asks how to increase it: shorter approval paths, fewer sign-offs, decision deadlines. I did the same in *The Tautai Principle*, which set a decision-latency benchmark of "< 72 hrs" and called latency a "direct proxy for judgment efficiency"; Appendix A records the departure.³² Chapter 1 gave the reason a fast loop is not in itself a good one: a fast loop on a wrong picture takes the organization further in the wrong

³²Mathis, K. (2025). *The Tautai Principle*, p. 69 ("< 72 hrs") and p. 201 ("direct proxy for judgment efficiency"); also p. 78 n. 40 and p. 80 on speed. Recorded as an edition trigger; Appendix A.

direction, faster. This chapter adds the mechanism. Speed is a consequence of model quality. Where the model is good, recognition is fast without being asked to be. Where the model is poor, the honest process is slow, because it has to compare, simulate and probe, and a speed target does not make the model better. It makes the organization act on the poor model sooner, and with the retrospective confidence that Weick's jurors show.

What speeds up a slow organization, on this reading, is investment in the model: experience formed under feedback, forums in which accounts can be compared, records that let the next round learn from the last, and the removal of option comparison from the places where recognition would do. Latency is worth measuring as an outcome. It tells the organization something about the state of its models. As a target it invites exactly the behavior it was meant to cure.

Chapter 1 drew the loop with an inner and an outer circle. A latency figure measures the inner one, from cue to action. What makes the inner circle fast is the outer one: the slow turn from action to learning to changed capability, which is where the commanders' models were formed and where the organization's are. An organization that shortens its approval paths while starving the outer circle gets, for a while, faster decisions on models that are no longer improving, and then the retrospective redefinition that turns each of them into a success. Chapter 11 takes up the outer circle. The commanders had spent years in it before they spent a minute in the inner one.

When a machine proposes the action

A model that ranks options is now cheap to build, and organizations are installing them at the point in the loop this chapter is about: a system that proposes the next best action to a service agent, that scores a list of initiatives, that drafts the recommendation before the meeting. Chapter 7 will make Hutchins' point that a tool changes the task rather than amplifying a capability. Applied here, a ranking model changes what the human selector does. The

field of possible actions is no longer held back and released by the person; it arrives pre-sorted, with one candidate at the top, and the person's task becomes the release or inhibition of a proposal.

Where the release sits is then the question. If the person can evaluate the proposal, the arrangement is a functional system of the kind chapter 7 describes, with the release still in human hands. If the person cannot, the release has moved to the model without anyone deciding that it should, and the organization has a selector whose grounds it cannot inspect. The belief core's rule, that delegation is legitimate only as deep as the capacity to evaluate what comes back, applies to selection with particular force, because a released action changes the terrain.³³ Chapter 13 takes this up with the rest of the machine thread.

What would count against it

The chapter has claimed that for experienced people understanding contains the choice, that selection and decision are different things needing different remedies, that decision rights belong where the model sits, and that speed is a consequence of model quality. Two findings would count against the core of this. The first is cycle speed predicting adaptive success independently of the quality of orientation, in comparable complex settings; if fast loops on poor models did as well as slow loops on good ones, the argument of the speed section would fail. The second is experienced practitioners in complex settings improving their outcomes by systematic option comparison rather than by better recognition; if they did, Klein's finding would not carry over from the fireground to the

³³The Tautai company belief core, on delegation and the capacity to evaluate; the wording "only as deep as the capacity to evaluate what comes back". Standalone readers: Appendix A.

organization. Where to look is the naturalistic decision-making literature collected by Zsombok and Klein, and Klein's own statements of the limits of recognition.^{34,35}

Part II has described the loop, from cue to account to action, as if one actor ran it. The commanders make that simplification look natural, because the recognition sat in one head and had been formed there over years. The *Palau* shows that it need not: a workable action was selected by a team that had no decider, held in a division of labor that nobody had represented. Part III drops the simplification. Its first question is where the model, and with it the choice, is actually held, and chapter 7 begins with a single bearing on its way from an instrument to a chart.

References

- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18(4), 513–549.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Mathis, K. (2025). *The Tautai Principle*.
- Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The structure of “unstructured” decision processes. *Administrative Science Quarterly*, 21(2), 246–275.
- Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge.
- Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76.

³⁴Zsombok, C. E., & Klein, G. (Eds.) (1997). *Naturalistic Decision Making*. Erlbaum.

³⁵Klein (1998), p. 24 (“fully 80 percent” recognitional; the novice case at the oil tank farm; comparison for novices, for justification and for novel situations).

- Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.
- Zsombok, C. E., & Klein, G. (Eds.) (1997). *Naturalistic Decision Making*. Erlbaum.

Chapter 7

Chapter 7 — Where understanding is held

A bearing on its way

On the bridge of the *Palau*, a pelorus operator sights a landmark on the shore and reads a bearing off the gyro-compass scale. He speaks it. The bearing recorder, who is not looking at the shore, receives it as “strings of spoken digits” and writes it into the bearing log.¹ The plotter, who has neither seen the landmark nor heard the digits, takes the numbers from the log and turns them into lines on the chart. Where the lines from several landmarks cross, the ship is. Chapter 1 gave this path in one sentence. It deserves a chapter, because of what is missing from it.

What is missing is a person who knows where the ship is. The pelorus operator holds a number. The recorder holds a row in a log. The plotter holds lines. Each of them could tell you what he has in front of him, and none of them, at the moment the fix is made, holds the ship’s position as a thought. The position is produced by the path the bearing travels, through three people and three media, and it exists on the chart before it exists in anyone’s head. When the navigator looks at the chart and says where the ship is, he is reading a result that the system has already computed.

Edwin Hutchins studied that bridge, and *Cognition in the Wild* (1995) is his account of what he found. Its central

¹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 4, p. 226 (“strings of spoken digits”; novices and the structure available in social relations).

move is stated in the introduction: he moves “the boundaries of the cognitive unit of analysis out beyond the skin of the individual person” and treats “the navigation team as a cognitive and computational system”.² The unit that thinks, in his analysis, is the functional system: the people, the instruments, the charts and logs, the procedures and the roles, together. Part II described the loop as if one actor ran it, and the commanders of chapter 6 made that look natural. Here the simplification is dropped, and the bearing is the reason.

The other theories in this book place understanding elsewhere. For Endsley, situation awareness is in the operator’s head, and displays and instruments are inputs to it. For Weick, meaning is made between people, in talk and action, and the artifacts are props. Hutchins places it in neither. The bearing is not in a head and it is not a meaning made in conversation; it is a state that moves through a system and is transformed at each step. Recent work on distributed team cognition, in the volume edited by Michael McNeese, Eduardo Salas and Endsley herself, builds the bridge from team situation awareness to this view.³ Orestis Palermos states the general thesis of the field: appropriately interacting individuals can form “collective entities with cognitive properties, over and above the sum of their individual members’ cognitive properties”.⁴

The path a representation travels

The claim that a team thinks would be a metaphor if it came without a mechanism, and Hutchins supplies one.

²Hutchins (1995), introduction (chapter summaries): unit of analysis; computation as propagation of representational state; local functional system; accumulation of structure in tools as a cognitive process; “two or more heads”.

³McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.

⁴Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 5, p. 115). CRC Press.

Chapter 3 of his book “elaborates a conception of computation as the propagation of representational state across a variety of media”, which “permits the use of a single language of description to cover cognitive and computational processes that lie inside and outside the heads of the practitioners”.⁵ The bearing is a representational state. It is held first in the alignment of a sight with a landmark, then in a number on a scale, then in sound, then in ink, then in a pencil line. Each medium holds it differently, and each transformation is done by somebody or something. Cognition, in this account, is what happens along that path.

The account matters for organizations because it makes their thinking observable. Where the ship’s position is held becomes an empirical question, and so do three others that this book asks of every organization it looks at. Where is this computation done? Which step is carried by a person, and which by an artifact? And where, along the path, does the picture get lost? Each step in the bearing’s journey can fail: a digit misheard, a row written in the wrong column, a line drawn from the wrong scale. Each failure has a location, and a system designed with the path in mind can put an error check where the failure is likely.

Weick, whose sensemaking chapter 4 developed, explains how meaning is constructed. Hutchins traces the path on which it travels. The two are not rivals; they answer different questions, and an organization needs both answers. It needs to know how its people build an account of a broken situation, and it needs to know how that account, once built, gets from the people who built it to the people who will act on it, and what happens to it on the way.

Consider the organizational equivalent of the bearing. A customer says something to a service agent. The agent

⁵Hutchins (1995), introduction (chapter summaries): unit of analysis; computation as propagation of representational state; local functional system; accumulation of structure in tools as a cognitive process; “two or more heads”.

files a ticket and picks a category from a list. The category is counted; the count appears on a dashboard; the dashboard number becomes a bullet in a monthly report; the bullet becomes a line in a slide for the board. At each step the state is re-represented in a new medium, and at each step something is dropped: the customer's words at the first, the agent's judgment about which category fitted at the second, the trend behind the number at the third. The organization's picture of its customers is the product of that path, in the same way that the ship's position is the product of the bearing's. Nobody in the organization holds the picture whole, and the picture that reaches the board has been computed by a system whose steps nobody designed as a whole. The three questions apply. Where along this path is the customer's meaning actually understood, and by whom? Which steps are done by people and which by the ticketing system's category list? And where is the picture lost? The place to look first is the second step, where a person's understanding is forced into a category that an artifact provides, and the artifact's categories were fixed when the ticketing system was configured.

Episode — From the operator's call to the board slide

"It keeps restarting since the building system was connected." The operator of an office building said this on the phone to Jana, a first-line agent in the after-sales service of Halvern, a manufacturer of heat pumps for commercial buildings, in the spring after the company's C4 controller generation had gone on sale. Jana had heard nearly the same sentence from two other operators that month. She opened a ticket, and the form asked her for a cause. It offered five: product defect, installation error, operating error, wear, other. The unit was doing what its controller told it to do. The last people to have changed anything were the installer partner's, who had connected the C4 to the building-management system through a gateway. She picked installation error and typed the operator's sentence into the free-text field.

Tobias, a field technician in region South, went to the site. He watched the compressor start and stop far too often whenever the building system overrode the set-point, wrote that into the free text as well, and closed the ticket under the category it already had. After that no person handled the sentence again. The free-text field was not carried into any report. Quality assigned the ticket a root-cause code, and because the five codes mapped one-to-one onto the five categories, the step that looked like a second opinion could only repeat the first.

The service dashboard showed ticket counts per category, per region, per week. In region South, installation error stood at about twenty a week in spring and at forty-seven by the end of the summer. Of roughly 11,000 C4 units installed that year, about 1,900 were connected to a building-management system, and the rising tickets came from those. The dashboard had no column for that, so a rise confined to one sixth of the units appeared as a rise among installers in general. Henrik, the Head of Quality, owned the template of the monthly quality report, which allowed one bullet per category with a trend arrow and a measure. The bullet read: installation errors, arrow up, measure: partner training. The Product Quality Board saw the bullet. In November a training program for installer partners was launched at a cost of about €180,000. The slide for the executive board that quarter carried one line on service: "Service cost up 9% owing to installation quality at partners; training program launched."

Every step after the first was carried out correctly. The count was accurate, the arrow pointed the right way, and the nine percent was the real figure. The operator's sentence had named a condition, the connection to the building system, and the form had no category for something that happened between two systems. What was lost when Jana chose one of five was the condition, and nothing further along the path could

put it back. The operator, Jana and Tobias had the sentence and never saw the report. The people who read the report and the slide had never heard the sentence, and no role at Halvern had both in front of it.

The path also shows what makes a system of this kind robust, and the *Palau* is again the case. When the gyrocompass failed, the team reconfigured itself and kept fixing the ship's position. Hutchins locates the reason in the overlap of knowledge between roles: it is "an important source of the robustness of such systems in the face of error and interruption". Because the members' knowledge overlaps, the team can "reconfigure dynamically", and "the individuals are a sort of flexible tissue that moves to ensure the propagation of task-relevant representational state".⁶ The image is exact. What the tissue protects is the path, and the path is what produces the position. An organization whose roles do not overlap has no tissue; when a step fails, the representation stops there, and the picture is lost at that point until someone notices it is missing.

Error correction works the same way, and here I extend Hutchins' argument rather than report his example: the overlap that lets the team reconfigure also lets it catch a misheard digit, because someone else is in a position to hear it and knows what is plausible. Hutchins draws a further consequence from error correction: "every error-correction event is a learning context not just for the person who commits the error but for all who witness it."⁷ Whether errors are witnessed is a property of how the work is arranged, and it decides whether the system learns from them or only recovers from them. Chapter 12 returns to this.

⁶Hutchins (1995), ch. 4, pp. 221 and 227.

⁷Hutchins (1995), ch. 6, p. 279.

What a tool does to a task

Hutchins, in his chapter on the tools of navigation, quotes a common view: “It has now become commonplace to speak of technology ... as an amplifier of cognitive abilities.”⁸ A calculator amplifies arithmetic, a chart amplifies spatial reasoning, a dashboard amplifies oversight. He then shows, following Michael Cole and Peg Griffin, that the appearance of amplification comes from looking at the product of the work rather than the process. Judged by the product, the navigator with a chart finds the position faster and more reliably than a navigator without one. Judged by the process, nothing has been amplified. “None of the component cognitive abilities has been amplified by the use of any of the tools. Rather, each tool presents the task to the user as a different sort of cognitive problem requiring a different set of cognitive abilities.”⁹ The chart does not make the navigator better at trigonometry. It turns a problem of trigonometry into a problem of drawing straight lines, which is a different problem needing different skills, and the trigonometry is done by the chart’s construction, once, by whoever made it.

Chapter 1 gave the sharpest form of this in the comparison between Western navigation and the navigation of the Caroline Islands, where the same task is done with an entirely different representation and therefore as a different cognitive problem. The general form is in Hutchins’ introduction: a person working with a tool forms a “local functional system” with “cognitive properties that are radically different from the cognitive properties of the person alone”.¹⁰ And the tools are made by the people who use

⁸Hutchins (1995), ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

⁹Hutchins (1995), ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

¹⁰Hutchins (1995), introduction (chapter summaries): unit of analysis; computation as propagation of representational state; local functional system; accumulation of structure in tools as a cognitive process; “two or more heads”.

them, or by their predecessors. “Humans create their cognitive powers by creating the environments in which they exercise those powers.”¹¹

This changes what a manager is doing when a tool is introduced. The usual account says the tool makes the people better at what they already do. Hutchins’ account says the tool changes what they do, and that the new task may call for abilities the people do not have, or make redundant abilities they have spent years forming. A forecasting model given to a planning team does not amplify the team’s judgment about demand. It changes the team’s task from forming a judgment to evaluating a forecast, and evaluating a forecast is a different problem, one that presupposes a model of the market good enough to say when the forecast is wrong. A dashboard given to a management team does not amplify oversight. It replaces the task of asking what is going on with the task of reading what the dashboard’s designers decided could be shown, and chapter 1 already said that a measurement system encodes an old picture. In each case the question to ask is the one Hutchins asks of the chart: what is the new cognitive problem, and who is now solving the old one, and where?

Solutions of earlier generations

The chart also shows something about time. The trigonometry it embodies was worked out long before the navigator was born, by people he will never meet, and it is available to him without effort. Hutchins describes the implementations of tasks as “part of a cultural process that tends to collect representations that permit tasks to be performed by means of simple cognitive processes”,¹² and he counts the accumulation of structure in the tools of a trade as “itself a cognitive process”.¹³ The formula he

¹¹Hutchins (1995), ch. 3, p. 169.

¹²Hutchins (1995), ch. 3, p. 155. The reading of the chart as a representation that turns trigonometry into line-drawing is mine, drawn from ch. 3 as a whole.

¹³Hutchins (1995), introduction (chapter summaries): unit of analysis; computation as propagation of representational state; local functional

gives at the end of the book puts culture and cognition on one footing: “a major component of culture is a cognitive process ... and cognition is a cultural process”.¹⁴

An organization, on this view, thinks in part with the solutions of earlier generations. Its forms, procedures, templates, category lists and reporting lines are representations that permit today’s tasks to be done by simple steps, because the hard thinking was done once and built in. That is a capacity: nobody on the bridge has to invent the chart, and nobody in a claims department has to invent the claims form. It is also an inertia, and the two cannot be had separately. A form that solved a problem well encodes the problem as it was understood when the form was designed; when the problem changes, the form goes on presenting the old one, and it presents it so smoothly that the people using the form do not experience it as a picture at all. The category list in the ticketing system is such a form. So is a strategy template, a risk register, an approval workflow.

The *Palau* provides the other half of the same lesson, and chapter 1 already told it: the team found a new way of organizing its work under pressure, and afterwards the solution “was not saved in the system”.¹⁵ Hutchins draws a general conclusion from the episode that this book takes as one of its foundations: “The processes by which work is accomplished, by which people are transformed from novices into experts, and by which work practices evolve are all the same processes.”¹⁶ Doing the work, learning to do it, and changing how it is done are one process seen at three time scales, and all three run through the same artifacts. What was not written into a procedure, a checklist or a role description after the *Palau* anchored was lost, however well it had worked. What is written into them stays, however badly it now fits. Chapter 11 takes up both

system; accumulation of structure in tools as a cognitive process; “two or more heads”.

¹⁴Hutchins (1995), ch. 9, p. 354 (the running head is illegible in the scan; the chapter’s offset places the formula on printed p. 354).

¹⁵Hutchins (1995), ch. 8, p. 351.

¹⁶Hutchins (1995), ch. 8, p. 351.

halves as organizational learning and unlearning. What this chapter adds is where they happen: in the artifacts as much as in the people.

Memory that nobody holds

If the picture is produced along a path, the organization's memory of its situation is held along a path too. Chapter 1 said that an organization's memory sits partly in people and partly in documents, tickets, budgets and the knowledge of who knows whom. Two literatures give this more shape, and the book has not yet verified either in the text, so they are stated with care.

James Walsh and Gerardo Ungson's account of organizational memory treats it as distributed across the organization rather than held in any archive, with people as only one of the places it is retained.¹⁷ Daniel Wegner's concept of transactive memory describes how members of a group come to rely on each other as external memory: each person remembers some things and remembers who remembers the rest, so that the group's memory is larger than any member's and depends on the members knowing each other's specialties.¹⁸ Kyle Lewis carried the concept into work teams.¹⁹ Hutchins himself asked, in a paper whose title is its argument, how a cockpit remembers its speeds.²⁰

¹⁷Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91. p. 62, p. 81 (six retention facilities).

¹⁸Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior*. Springer. p. 186 ("A transactive memory system is a set of individual memory systems in combination with the communication that takes place between individuals"), p. 190 ("Other people can be locations of external storage"), p. 191 (learning each other's domains of expertise).

¹⁹Lewis, K. (2003). Measuring transactive memory systems in the field. *Journal of Applied Psychology*, 88(4), 587–604. p. 587 (field measure of transactive memory systems in work teams; 15-item scale).

²⁰Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19(3), 265–288. p. 283 ("Speed bugs do not help pilots remember

The practical form of transactive memory is familiar to anyone who has worked in an organization for more than a year. There is a person everyone asks about the legacy system; a person who knows why the contract with the distributor has that clause; a person who remembers what was tried in 2019 and why it failed. The organization's memory of those matters is the fact that people know to ask those people. When one of them leaves, what is lost is the pointer as well as the knowledge, and the organization discovers the loss only when a question arrives that nobody can now route. The same holds for artifacts. A wiki that nobody knows to consult is not memory; the ticketing system's history is memory only for the people who know how to query it and that it is worth querying.

For this chapter the consequence is that memory, like the position of the ship, is a property of the path and not of any store. Where an organization keeps its account of what it once expected, which is chapter 6's decision record, and who knows to look there, is a question of design.

An operating model is a thinking model

Chapter 1 stated in a paragraph what this chapter can now say with the mechanism in hand. An operating model, the arrangement of roles, forums, artifacts, handoffs and reporting lines by which an organization does its work, is also the arrangement by which it thinks, because the thinking is the propagation of representations through exactly that arrangement. Three design questions follow, each of which has an answer in every organization whether or not anyone has asked it.

Where does working memory sit? On the bridge, the log and the chart hold the state of the work between one fix and the next, so that no person has to. Between the meetings of a management team, what holds the state of its understanding? If the answer is the memory of the people who were in the room, the working memory of the

speeds; rather, they are part of the process by which the cockpit system remembers speeds.”), p. 286 (“The cockpit system remembers its speeds”).

team is as good as the attendance at the next meeting. If the answer is a document that the next meeting reads first, the team has an artifact doing what the bearing log does.

What detects errors? The bridge, as said above, catches a misheard digit because of overlap. In an organization, error detection is a matter of who sees a representation before it is acted on and whether they have a model good enough to see that it is wrong. A forecast that goes from the analyst to the decision without passing anyone who could doubt it has no error check, whatever the approval process says. Hutchins notes that novices can take part in the work before they understand it, because “much of the structure required for them to organize their activity is available in the social relations”,²¹ and that what people can learn depends on their horizon of observation, which tools can widen or narrow: “the design of a tool may change the horizon of observation of those in the vicinity of the tool”.²² A bridge on which everyone can hear the bearings called out has a wide horizon; a system in which each person sees only their own screen has a narrow one, and error detection narrows with it.

How does meaning cross handoffs? Every boundary in the operating model is a place where a representation is re-coded for a new medium and a new reader. Beer’s Viable System Model describes these boundaries as channels and transducers and asks whether they can carry the variety that has to cross them; chapter 9 develops that. Hutchins shows what happens inside a transducer: a state in one medium becomes a state in another, and something is preserved and something is dropped. The handoff from the agent’s understanding to the ticket category is a transducer, and so is the handoff from the unit’s account of its market to the line in the corporate plan. Whether the meaning that crosses is the meaning that was sent is not a matter of communication skills. It depends on what the medium on the far side can represent.

²¹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 4, p. 226 (“strings of spoken digits”; novices and the structure available in social relations).

²²Hutchins (1995), ch. 6, p. 270.

Representations as organization design

The three questions have a common answer: the representations. Which categories the ticketing system offers, what a decision record must contain, how a metric is defined and who may redefine it, what a problem statement looks like before a team is allowed to start: these decide where working memory sits, what errors can be seen, and what survives a handoff. Sensemaking Article 3 calls the deliberate management of these things representational governance and treats decision records as the organization's memory of what it committed to and why.²³ Four kinds of representation carry most of the weight.

Problem definitions fix what a team is looking at and therefore what it can notice; a problem defined as "reduce ticket volume" and one defined as "find out why customers call" send the same team down different paths. Decision records, from chapter 6, hold the hypothesis a commitment rests on so that the next round of understanding can test it. Metric dictionaries fix what a number means across the units that report it, and an organization without one has units that agree on a figure and disagree on what it measures. And boundary objects, in the term Susan Leigh Star and James Griesemer introduced, are representations that several groups use in common while each reads them in its own terms, which is what a roadmap, a budget line or a shared backlog does between functions.²⁴ Each of these is usually filed under communication: a matter of presenting understanding that already exists. On Hutchins' account they are part of the computation. The representation is where a piece of the organization's thinking is done, and changing it changes the thinking, as changing the chart would change how the ship's position is found.

²³Sensemaking Article 3, 51-Stories/Sensemaking Articles/article-03-representational-governance.md, "Decision memory (decision records)".

²⁴Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, "translations" and boundary objects. *Social Studies of Science*, 19(3), 387–420. p. 393 (boundary objects "plastic enough to adapt to local needs ... robust enough to maintain a common identity across sites").

There is a further reason to treat them as design objects. Experts, the distributed-cognition literature on expertise observes, distribute their cognition into their context and their artifacts; novices cannot yet do this and hold everything in the head.²⁵ An organization that designs its representations well gives its experienced people something to think with and its novices a structure to take part in before they understand it. One that leaves its representations to accumulate gives both of them a path whose losses nobody has mapped.

A new participant

Machines have always been in the functional system; the gyrocompass was one, and the *Palau's* team found out what it had been doing when it stopped. What has changed is that a model can now sit at any point in the chain of representations and do a step that a person did. It can read the customer's words and pick the category. It can take the dashboard's numbers and write the bullet for the report. It can take the bearing log and draw the line.

Hutchins' argument about tools applies to this participant without modification. The model does not amplify the person whose step it takes; it changes the task of the people on either side of it. The person before it is now producing input for a model rather than a message for a colleague, and the person after it is evaluating a model's output rather than reading a colleague's. The chapter's three questions then apply with unusual force. Where in the chain does the model sit, and which representation does it consume and produce? Can its outputs be inspected, so that the error check that overlap provided among people has an equivalent here? And what has happened to the horizon of observation: can the people around the model see what it is doing, in the way the bridge could hear the

²⁵Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 8, pp. 187–191). CRC Press.

bearings, or has a step of the organization's thinking become invisible to the organization? Chapter 13 takes these up with the rest of the machine thread. Here it is enough to note that a functional system with a model in it is a new system with properties of its own, some of which nobody yet knows, and that the way to find out is to trace the path.

What the lens does not see

Hutchins studied well-structured work. Where the ship is, is a matter of computation; the task is defined, the representations are standardized, and the crew does not have to ask what a bearing means. The strength of the lens comes from this, and so does its limit. Surprise, equivocality and the collapse of a frame, which chapter 4 treated, do not appear on the bridge, and the propagation of representations says nothing about what happens when nobody knows what should be represented. Rank and military identity appear in Hutchins' account, but who decides which account counts is not his theme, and the politics of the picture remain with chapter 10. And distributed cognition describes how a system thinks, not what it is for; viability and purpose stay with Beer in chapter 9.

The chapter also carries a disagreement into Part IV, and it should be stated rather than smoothed. Ikujiro Nonaka, whose theory of knowledge creation organizes chapters 11 to 14, holds that knowledge "is created by individuals" and that the organization "'organizationally' amplifies" it.²⁶ For Hutchins the system is what thinks, and the *Palau* is his evidence: a solution found by the organization before any member had it. One reading that keeps both is that Nonaka describes how an individual's contribution becomes shared and Hutchins describes where the shared

²⁶Nonaka (1994), p. 17 (individuals create knowledge; the organization amplifies it; intention, autonomy and fluctuation; Weick 1979 on enactment), p. 18 (breakdowns after Winograd and Flores). The gloss of "breakdown" as the moment a routine stops working and its assumptions become visible is the author's; Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

result then lives and how it works. The book does not settle the matter here. Where it uses Nonaka's vocabulary in Part IV, it will say what the Hutchins reading would add.

What would count against the chapter's claim is simple to state. If changing an organization's representations, artifacts and handoffs left its coordination, error detection and learning unchanged, then the chain of representations would add nothing to what a description of roles and information flows already gave, and the chapter's design consequences would be empty. Where to look is the transactive-memory research of Lewis, and the research in human-computer interaction and computer-supported cooperative work on shared artifacts, which Jim Hollan, Hutchins and David Kirsh set out as a foundation for human-computer interaction research.²⁷

A system of people and artifacts can hold an understanding that no member holds, and the bearing shows how. Hutchins is careful about what follows from this, and the next chapter takes up his caution: "It is not the case that two or more heads are always better than one."²⁸ Whether the system understands more than its members or less depends on how it is organized, and that is chapter 8.

References

- Cole, M., & Griffin, P. (1980). Cited in Hutchins (1995), ch. 3.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196.

²⁷Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196. The paper sets out distributed cognition as a foundation for human-computer interaction generally, not a program on shared artifacts specifically.

²⁸Hutchins (1995), introduction (chapter summaries): unit of analysis; computation as propagation of representational state; local functional system; accumulation of structure in tools as a cognitive process; "two or more heads".

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19(3), 265–288.
- Lewis, K. (2003). Measuring transactive memory systems in the field. *Journal of Applied Psychology*, 88(4), 587–604.
- McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In McNeese, Salas & Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*, ch. 5. CRC Press.
- Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In McNeese, Salas & Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*, ch. 8. CRC Press.
- Sensemaking Article 3, “Decision memory (decision records)”.
- Star, S. L., & Griesemer, J. R. (1989). Institutional ecology, “translations” and boundary objects. *Social Studies of Science*, 19(3), 387–420.
- Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91.
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior*. Springer.

Chapter 8

Chapter 8 — When groups understand more, and when less

Two or more heads

“It is not the case that two or more heads are always better than one.”¹ Edwin Hutchins put that sentence in the introduction to *Cognition in the Wild*, a book whose whole argument is that the unit that thinks is larger than one head. He had watched a navigation team do together what none of its members could have done alone, and he had traced, step by step, how the team’s picture of where the ship was passed from instrument to voice to log to chart (chapters 1 and 7). He still put the warning up front. A system of people and artifacts can understand more than any of them. It can also understand less, and the same book shows how.

Chapter 7 placed understanding in such systems and followed a representation through one. This chapter asks what a system gains, or loses, by being a group: a team, a leadership circle, a network of experts, a committee. The answer depends on conditions, and the conditions can be named. They concern who is in the group, whether its members will say what they see, and, more than the literature on teams usually admits, how their communication is organized.

Chapter 1 said that each of the five activities of the loop has a collective form, and the book has carried that as a

¹Hutchins (1995), introduction (chapter summaries; the local functional system).

thread: who shares a cue, who contests an account, who commits. This chapter is the thread's home. It is easy to say that perceiving, understanding and choosing are done together; it is harder to say when doing them together is an improvement. A group can notice more than one person, because its members stand in different places; it can also notice less, because each member assumes the others are watching. It can build a richer account of a surprise, or it can settle on the first account faster than any member would have alone. The question is the same for every activity of the loop, and so is the shape of the answer.

Conditions of collective intelligence

In 2010 Anita Woolley and her colleagues reported an experiment that has shaped the discussion since. They gave small groups a series of unrelated tasks and asked whether a group that did well on one did well on the others, as an individual does across the sub-tests of an intelligence test. They reported that it did: a single factor, which they called collective intelligence, accounted for much of a group's performance across tasks. What predicted that factor was the members' social sensitivity, how evenly they took turns in conversation, and the composition of the group, far more than the average intelligence of the members or the intelligence of the cleverest of them.²

The finding has traveled far, and it matches what anyone who sits in meetings for a living has seen: a group of very capable people producing less than one of them would alone, and a group of ordinary people producing something none of them could have. Whether the factor holds has been argued over since. Credé and Howardson re-examined the evidence in 2017 and questioned how

²Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688. p. 686: the c-factor is correlated with members' social sensitivity, equality of conversational turn-taking and the proportion of females, not with the average or the maximum individual intelligence; pp. 687–688.

much of it supports a general factor; Riedl and colleagues returned to the question in 2021 with a larger body of data.³ I do not build on a settled finding here, because there is none. What I take from the debate is narrower. Whether a group understands well is a property of the group, it can be studied as such, and what predicts it is how the members interact, more than how clever they are one by one.

The conditions that recur across this literature are four. Diverse perspectives, so that the group holds more than one account of the situation. Equal participation, so that the accounts get said. Social sensitivity, so that members notice what the others know and feel. Structured interaction, so that the group does something with its differences instead of averaging them away.⁴ Each is a condition of the group's understanding; none is a condition of its harmony.

Read against the loop, the four do different work. Diverse perspectives serve perceiving: a group whose members face different parts of the environment holds more cues than any of them, provided the cues are brought in. Equal participation and social sensitivity serve understanding: an account can only be built from what is said, and it can only be corrected by someone who notices that a colleague has gone quiet. Structured interaction serves the step from several accounts to one, which is where a group most often loses what it had gained. A group that meets the first three conditions and lacks the fourth is a group of well-informed, attentive, articulate people who talk their differences away in the first ten minutes. The

³Credé, M., & Howardson, G. (2017). The structure of group task performance: A second look at "collective intelligence". *Journal of Applied Psychology*, 102(10), 1483–1492. p. 1483: re-analysis of six samples finds "the empirical support for the construct is generally weak ... insufficient support for the existence of a collective intelligence construct". Riedl, C., Kim, Y. J., Gupta, P., Malone, T. W., & Woolley, A. W. (2021). Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21). Supports the c-factor: 22 studies report "a robust CI factor" (not a critique).

⁴The four conditions as listed in the outline for this chapter; they summarize the collective-intelligence and team-learning literature and are not taken from a single source.

fourth condition is the one organizations neglect most, and Hutchins explains why it matters.

More communication is not always better

Hutchins asks a question the collective-intelligence literature mostly leaves alone: what happens to a known property of individual minds when minds are connected? The property he chose is confirmation bias, “a propensity to affirm prior interpretations and to discount, ignore, or reinterpret evidence that runs counter to an already-formed interpretation”.⁵ Everyone has it. Common sense says a group corrects for it, because one person’s early interpretation meets another’s. Hutchins tested the common sense with simulated groups of interpreters, varying how richly their members communicated. The simulations indicate “that more communication is not always in principle better than less. Under some conditions, increasing the richness of communication may result in undesirable properties at the group level.”⁶

My reading of what those conditions are, which is a reading of mine and does not describe the simulation’s design, is this. A group whose members exchange interpretations constantly from the start converges on whichever interpretation formed first, and then holds it together: the individual bias has become a group bias, with more weight behind it. A group whose members work on the evidence apart for a while and compare interpretations later holds several accounts open long enough to test them against each other. On this reading a group understands more than its members only when its members’ interpretations are allowed to differ before they converge.⁷

The organizational version needs no simulation. A leadership team that meets every morning, shares one dashboard and prides itself on alignment has organized

⁵Hutchins (1995), ch. 5, p. 239.

⁶Hutchins (1995), ch. 5, pp. 239 and 252.

⁷Author’s reading of Hutchins (1995), ch. 5, pp. 239–252. The design of the simulations and the mechanism (interpretations are allowed to differ before they converge) are Hutchins’ own at pp. 255–256.

its communication so that the first plausible account of a surprise becomes the shared account within hours, before anyone has looked at the evidence alone. Alignment, which the team counts as a strength, is the group-level property the simulation warns of. The more the members trust each other, the faster it happens. The remedy is structural and cheap: separate the time in which accounts are formed from the time in which they are compared, and let people form them apart.

Episode — The January week

“It’s the installers. The cold is only showing us what was already there.” Henrik, Head of Quality at Halvern, said it on the second morning of a cold week in January, about ten minutes into the daily meeting of the service leadership team. In the heating season the team met every morning at a quarter to eight: Katrin as Head of Service, Henrik, and the leads of dispatch, parts, the support desk and the second line. They stood in front of the one service dashboard, which counted tickets per category, per region, per week. Compressors in the C4 controller generation had begun to fail over the weekend, and the tallest bar on the screen was installation error.

The account fitted everything the room had in front of it. Installation error had climbed since the summer. The monthly quality report had said so, a partner training program costing about €180,000 had been running since November, and the board slide named installation quality at partners as the reason service cost was up 9%. Low outside temperatures load a compressor harder, so a badly installed system could be expected to fail in a cold week. Katrin asked whether anyone read it differently. The dispatch lead did not. The second-line lead said the escalations he had seen were coded as installation cases, which they were. Nobody at the table had come with another account, because

nobody had looked at the failed units before the meeting. Each of them had looked at the dashboard, and they had looked at it together.

By nine the work was distributed. Parts ordered compressors, dispatch moved the affected sites to the top of the list, and Katrin phoned the two largest installer partners to say their crews would be called to training earlier. The team was known in the company for answering with one voice, and Katrin counted on that. Before noon she sent Dr. Reuter, the COO, a three-line note that gave the cause as installer errors made visible by the cold. The account was then about four hours old.

It did not change during the week. By the sixth day 14 compressors had failed, and each morning's new failures were added to the count as further cases of the same thing. On the fourth morning one of the leads proposed a monitoring dashboard, so that the next cold spell would show sooner. The tickets for the failed units had free-text fields, and some of them said when the restarting had begun. They stayed unopened that week, and nobody sorted the 14 sites by what else the controller was connected to. Milan, the senior second-line specialist who had seen such cases in the summer, was not in the meeting.

The team held the account through the winter. Halvern replaced 96 compressors over the season, against about 25 in a normal one. The two large installer partners, blamed and sent to training, moved part of their business elsewhere. The account gave way only in March, when Katrin brought three individually traced cases to the Product Quality Board.

Overlapping knowledge

The same book gives the other side. Chapter 4 of *Cognition in the Wild* asks how the navigation team keeps working when a member is interrupted, makes an error or is

missing. Part of the answer is that the members' knowledge overlaps. Navigators come up through the simpler roles on the way to the harder ones (chapter 12 follows that trajectory), so each of them knows what the next person's work should look like and can see when it goes wrong.⁸ Hutchins calls this overlap "an important source of the robustness of such systems in the face of error and interruption". Because of it the team can "reconfigure dynamically", and, in his image, "the individuals are a sort of flexible tissue that moves to ensure the propagation of task-relevant representational state".⁹

That is the *Palau* of chapter 1 in one image: when the gyrocompass failed, the tissue moved, and the picture of where the ship was stayed alive. It is also what Mann Gulch lacked. A crew put together for one afternoon's job, whose members had not worked through each other's roles, had no overlap to move into when its frame collapsed.¹⁰

Overlap is redundancy, and redundancy is what efficiency programs remove first. A team in which every member knows only his own step is cheaper to run and easier to replace piece by piece, and it has no way to keep its understanding when a piece fails. Nonaka, from a different tradition, arrives at the same principle from the other direction. Among the conditions under which an organization creates knowledge he names redundancy, described, in the passage as it has been reported to me, as the conscious overlapping of information, activities and responsibilities across the company.¹¹ Hutchins' tissue is that principle seen at work on a bridge.

⁸Hutchins (1995), ch. 6, p. 264.

⁹Hutchins (1995), ch. 4, pp. 221 and 227.

¹⁰Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633: "Minimal organizations, such as we find in the crew at Mann Gulch, are susceptible to sudden losses of meaning"; the same page gives the cosmology-episode definition quoted in chapter 1, verbatim. Primary account: Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.

¹¹Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, about pp. 27–29 (organizational enabling conditions, redundancy).

Put the two findings side by side and the shape of an answer appears. A group understands more than its members when its knowledge overlaps and its interpretations are allowed to differ. It understands less when its knowledge is partitioned and its interpretations converge early. Both are properties of how the group is organized, and neither shows in the CVs of its members.

How much overlap a group needs is a design decision, and it differs by what the group is for. A team that operates something, a ward, a bridge, a trading desk, needs enough overlap that any member's error is visible to another and any member's absence can be covered, because its picture has to survive the failure of a part in real time. A team that decides something needs less overlap of skill and more overlap of exposure: its members should have seen the same environment from different places, so that their accounts differ in a way that can be compared. Chapter 14 takes both into its grids. The mistake to avoid at either level is to buy overlap where it is cheap, in meetings, and to remove it where it is dear, in people who have done each other's jobs.

Scale: trust in teams, interfaces in networks

Hutchins noticed that novices on the bridge could take part before they understood the task, because "much of the structure required for them to organize their activity is available in the social relations".¹² The people who hold each other's work in view stand within earshot. Their overlapping knowledge is partly knowledge of each other: who is reliable at what, who is new, whose numbers to check twice. It rests on relationships, and relationships do not scale. Robin Dunbar's proposal that there is a cognitive ceiling on the number of relationships a person can maintain is the usual reference here; I use it with care, because both the evidence for the ceiling and the figure that usually travels with it are debated.¹³

¹²Hutchins (1995), ch. 4, p. 226.

¹³Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6), 469–493; and the

What follows does not depend on the figure. Above the size at which everyone knows everyone, understanding cannot travel on trust, because there is nobody there to trust. It has to travel on interfaces: the representations of chapter 7, the forums of this chapter, the transducers of chapter 9. A network of teams has no flexible tissue. It has handoffs, and each handoff is a place where a representation is re-coded for the next medium and meaning is kept or lost. The two levels therefore need different things. A team needs overlap, and time to form interpretations apart. A network needs representations that carry an account across a boundary without the person who made it, and forums in which the accounts of different teams can be compared. Organizations regularly design one level as if it were the other: they hold alignment workshops for two hundred people where they need decision records, and ask a team of eight to talk to each other through tickets.

What a network keeps of the team's overlap is knowledge of a thinner kind: knowing who knows what. Chapter 7 introduced it under the name the psychologists gave it, transactive memory, and treated it as part of where an organization's understanding is held.¹⁴ Here it is the network's substitute for tissue. A team member can see that

later extension to human groups. The 1992 paper is about non-human primates and gives **no human figure**; the "150" comes from Dunbar (1993), not this paper. Dunbar, R. I. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences*, 16(4), 681-735. The 1993 paper carries the human figure: its reduced-major-axis regression of mean group size on neocortex ratio ($\log N = 0.093 + 3.389 \log CR$; accounting for 76% of the variance, $r^2 = 0.764$, over 36 primate genera) predicts, at the human neocortex ratio $CR = 4.1$, a group size for humans of 147.8 (95% CL 100.2-231.1), p. 682: "Equation (1) yields a predicted group size for humans of 147.8." The mean size of the intermediate-level groupings in the paper's ethnographic sample is 148.4 (range 90-221.5, $N = 9$), p. 685. Cited "with care" in the outline; the figure that circulates with the claim is deliberately not given here.

¹⁴Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior* (pp. 185-208). Springer. Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88(4), 587-604. Wegner 1987, pp. 186, 190, 191 ("Other people can be locations of external storage"; a transactive memory system "begins when individuals learn something about each

the next person's work is going wrong; a network member can at best know whom to ask. That is a weaker safeguard, and it fails in a characteristic way: the directory of who knows what is itself held in people's heads, is rarely written down, and leaves with them. An organization that reorganizes often loses less of its explicit knowledge than of this directory, and then wonders why its teams, each of which still knows what it knew, no longer understand each other.

There is a third level, which the book keeps in view without treating it here: the organization in its environment of other organizations, which chapter 9 reaches through recursion. At that level the interfaces are contracts, standards and the accounts each side holds of the other, and the question of this chapter, whether the whole understands more than its parts, is asked of an ecosystem rather than a team.

Forums that separate frame competition from commitment

Collective sensemaking, in Weick's sense (chapter 4), is what a group does when its picture breaks. Two things then have to happen, and they pull against each other. Several accounts of the situation have to be built and compared, which needs time, disagreement and a tolerance for half-formed stories. And one account has to be committed to, which needs closure, an owner and a record. A meeting that tries to do both at once does neither well. The competition of frames is cut short by whoever is senior or speaks first, and the commitment is never quite made, because it was made in the same breath as the argument and nobody wrote it down.

In an earlier article on sensemaking I proposed a *commitment forum* whose only output is a "decision record, owners, metrics, review date"; chapter 6 uses it as the place where a decision, in the book's sense of an explicit,

others' domains of expertise"); Lewis 2003, p. 587 (15-item field measure of transactive memory systems in work teams).

accountable commitment, is made.¹⁵ Here it does a second job. If commitment has its own forum, the meetings before it are free to be what Hutchins' simulation asks for: places where accounts formed apart are compared without the pressure to close. The separation also protects the record. Weick, following Garfinkel's study of jurors, observes that "the person defines retrospectively the decisions that have been made".¹⁶ A group does the same, more thoroughly, because it has more people to agree on the new version. Without a record, the group's account of what it decided drifts with the outcome, and there is nothing left for chapter 11's learning to test against.

In practice the separation is two meetings, or two halves of one, with a rule between them. In the first, people bring accounts they have formed apart, and the work is to make each account explicit enough to be compared: what cues it rests on, what it expects to happen next, what would show it to be wrong. Nobody decides. In the second, one account is committed to, an owner is named, the expectation is written down as the thing the next review will check, and the accounts that lost are kept, because one of them may be the right one next month. The rule between the two is that the first meeting ends without a decision even when everyone agrees, since the agreement of a group whose members have not yet looked at the evidence apart is what Hutchins' simulation warns of. Chapter 14 turns this into a design element; here the claim is only that the two kinds of work need two kinds of place.

Who sits in the commitment forum is the next question, and Beer answered it in *Brain of the Firm* in a passage that reads as if written for this chapter. Following Warren McCulloch he describes the "redundancy of potential command": "the command centre changes from time to time. Its location is a function of the information

¹⁵Sensemaking Article 2, 51-Stories/Sensemaking Articles/article-02-sensemaking-moves.md, "Blueprint C — Commitment forum". Article 3 on decision records as organizational memory.

¹⁶Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 10–11, on Garfinkel's jurors.

available”. Then the example: “Who, for example, really decides to buy ... It is the little group of employees, possibly very junior, which understands the need”.¹⁷ Decision rights sit where the understanding sits, and a forum that separates frames from commitment can move its commitment to that little group. Chapter 10 returns to what happens when the senior people would rather it did not.

Psychological safety: a precondition, not a mood

A group can understand only what its members say, or show. Amy Edmondson gave the condition its name in 1999: team psychological safety, a shared belief among members that the team is safe for interpersonal risk-taking, for asking, admitting error, disagreeing. In a field study of work teams she found it associated with learning behavior, and through it with performance.¹⁸ The finding has since been softened in use into a matter of atmosphere, something a team has if it feels good to be in. That is a misreading. A team can be pleasant and unsafe, because nobody there would risk the pleasantness; and it can be abrasive and safe, because everyone there knows an objection will be heard and will cost nothing. Safety is a condition of the group’s input. It decides which of the cues its members hold ever reach the group at all.

The condition scales badly with hierarchy. Chapter 2 treated organizational silence as one of the three filters on weak signals: people who see a cue and do not say it, because saying it would cost them something and the cost falls on them alone.¹⁹ In a group with a senior member

¹⁷Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

¹⁸Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383. p. 350: “team psychological safety—a shared belief held by members of a team that the team is safe for interpersonal risk taking”; a study of 51 work teams in a manufacturing company; “learning behavior mediates between team psychological safety and team performance”.

¹⁹Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725. p. 706 (organizational silence defined), p. 708 (a climate of silence: futile and/or dangerous).

present, the cost is highest for exactly the cues that matter most, the ones that contradict the account the senior member has already given. A leadership team is therefore the group in which safety is hardest to establish and most consequential when absent, and the group whose members most often report that it is fine. Edmondson's construct is a shared belief, which means it can be asked about, and the answers of the junior members count for more than the answers of the senior ones.

Hutchins gives the mechanism from the bridge: "every error-correction event is a learning context not just for the person who commits the error but for all who witness it."²⁰ An error corrected in the open teaches the whole team what the work should look like; an error that is hidden teaches nobody, and the tissue that would have moved to cover it does not know it is needed. Safety is therefore also the condition under which overlapping knowledge can do its work. And it is fragile in exactly the situations in which it matters. Mann Gulch is the extreme case: when the frame went, the role structure went with it, and the men who might have followed the foreman into the burnt grass had lost the structure in which his action could have made sense to them.²¹ Fear narrows what a group can build. The emotional side of collective sensemaking has its own literature, which I point to and have not worked through.²²

²⁰Hutchins (1995), ch. 6, p. 279.

²¹Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633: "Minimal organizations, such as we find in the crew at Mann Gulch, are susceptible to sudden losses of meaning"; the same page gives the cosmology-episode definition quoted in chapter 1, verbatim. Primary account: Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.

²²Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247. p. 222: emotion "signals the need for and provides the energy that fuels sensemaking" and affects the plausibility of accounts; p. 227: intense emotions "may forestall sensemaking...", moderate emotions energize it most.

Narrative: how accounts travel

Chapter 4 made the case that an account is a story: a plausible arrangement of cues, a frame that gives them a meaning, and an action the frame calls for, told in an order. The claim of this chapter is smaller and follows from it. Stories are the medium in which accounts travel between people, and a forum is a place where stories are compared. A dashboard carries the cues without the frame; a decision record carries the commitment without the story that led to it; a story carries all three, which is why it moves through an organization so much faster than either, and why it is so hard to correct once it has moved. Nonaka has a related observation from the side of knowledge creation: metaphor, he writes, “plays an important role in the externalization process”, the first step by which something tacit becomes sayable at all.²³ A story is how a member of a group gets a half-formed account out of her head and onto the table where the others can work on it. The organizational-narrative literature has developed this far beyond what I use here, and I point to it rather than summarize it.²⁴

Stories also carry time, which is why the book’s thread on temporality runs through them. Weick’s sensemaking is retrospective: the account is built after the fact, out of what has already happened (chapter 4). A story told in a forum is the group’s retrospect, and it fixes the order of events for everyone who hears it, including the order in which the cues were noticed, which is rarely the order in which they were there to be noticed. The forward-looking counterpart, the story of what will happen if we do this, is the form in which a group’s projection travels, and it is tested in the same way, by comparing it with another

²³Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

²⁴Boje, D. M. (1991). The storytelling organization: A study of story performance in an office-supply firm. *Administrative Science Quarterly*, 36(1), 106–126. p. 106: stories are “terse, requiring the hearer to fill in silently major chunks”; “organization as a collective storytelling system”. Czarniawska, B. (1997). *Narrating the Organization: Dramas of Institutional Identity*. University of Chicago Press.

group member's story of the same future. A group whose forums allow only the first kind of story explains well and anticipates poorly.

The design consequence is that the forums of this chapter have to make room for stories, and then discipline them. Room, because an account that has to be presented as a slide with three bullet points arrives without the cues that would let anyone test it. Discipline, because a story that is compared with no other story becomes the group's account by default, which is the confirmation-bias result once more, in narrative form.

Ba: the shared context

Nonaka, Toyama and Konno gave the shared context in which all this happens a name, *ba*, a Japanese word for place: the physical, virtual or mental space in which knowledge is created and shared. They distinguish four kinds, one for each of the four modes in Nonaka's theory of knowledge creation: an originating *ba* for the sharing of tacit knowledge face to face, a dialoguing *ba* for putting it into words, a systemizing *ba* for combining what is already explicit, and an exercising *ba* for making explicit knowledge one's own again in practice.²⁵ The paper defines the four kinds along two dimensions — whether the interaction is individual or collective, and whether it is face-to-face or through virtual media — and links each to a step in the conversion process, though it stresses that the links are “by no means exclusive”. The book's use of the concept is limited to what the four names carry, which is enough. What this chapter has called forums, and chapter 7 called representations, are attempts to build *ba* on purpose; the four kinds say that a group needs different places for different work, and that the systemizing kind,

²⁵Nonaka, I., Toyama, R., & Konno, N. (2000). SECI, *ba* and leadership: A unified model of dynamic knowledge creation. *Long Range Planning*, 33(1), 5–34. The four types of *ba* (originating, dialoguing, systemizing, exercising) are defined at p. 15 (definitions pp. 15–17), each linked to a conversion mode.

the only one most organizations build with any care, is the one in which the least is created.

The concept also brings the disagreement this chapter inherits from chapter 7 into the open. For Nonaka, knowledge “is created by individuals”, and the organization amplifies and shares what individuals have made.²⁶ For Hutchins, the system of people and artifacts is what thinks, and it can find a solution before any member has understood it: the *Palau*’s new division of labor was “discovered by the organization before it was discovered by any of the participants” (chapter 1).²⁷ Palermos states the general thesis of Hutchins’ side: appropriately interacting individuals can form “collective entities with cognitive properties, over and above the sum of their individual members’ cognitive properties”.²⁸ The two positions are in tension only if each is taken as the whole story. A reconciliation is available (Nonaka starts from the individual and describes how individual contributions become shared, Hutchins starts from the system and describes where the shared result lives and how it works), and the book leaves the question open until Part IV, where it has consequences for what learning is.

What the evidence does and does not support

The chapter has made four claims, of unequal strength. That whether a group understands well is a property of the group, predicted by how its members interact: supported by Woolley’s finding, contested in its generality,

²⁶Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 17 (knowledge created by individuals; intention, autonomy, fluctuation), p. 18 (the “breakdowns” passage after Winograd and Flores). (public PDF, pp. 14–21).

²⁷Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

²⁸Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 5, p. 115). CRC Press.

and held here in the weak form. That more communication can make a group understand less: supported by Hutchins' simulations and by the ordinary experience of aligned teams; the mechanism I gave is my reading, and the simulations should be read before it is relied on. That overlapping knowledge is what lets a group keep its picture when a part fails: supported by Hutchins' observation of the navigation team and by the contrast with Mann Gulch. That forums which separate the competition of frames from commitment, and the safety to say what one sees, are conditions of collective understanding: argued from the above, with Edmondson's study as the one empirical anchor.

What would count against these claims can be stated so that a reader can look. Groups that meet the conditions and still do no better than their best member on equivocal problems. The collective-intelligence factor failing to replicate, which would remove the first claim's support without touching the others. Groups with dense communication that show none of the group-level pathology the simulations predict. Forums that separate frames from commitment producing accounts no better than those of undifferentiated meetings. Where to look: Woolley against her critics; Edmondson's own work and what followed it; and Hutchins' fifth chapter.

What the chapter cannot answer alone is the question it keeps arriving at. Overlap, separate phases, safety, forums, *ba*: all of them are ways of organizing communication, and organizing communication is what structure does. Whether a cue formed in one team reaches another, how many layers it crosses on the way, what is kept and what is lost at each crossing, and at which level a picture of the environment is held at all: these are settled by roles, layers and channels before any group sits down. Chapter 9 takes the structure for what it is, a cognitive architecture.

References

- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.

- Boje, D. M. (1991). The storytelling organization: A study of story performance in an office-supply firm. *Administrative Science Quarterly*, 36(1), 106–126.
- Credé, M., & Howardson, G. (2017). The structure of group task performance: A second look at “collective intelligence”. *Journal of Applied Psychology*, 102(10), 1483–1492.
- Czarniawska, B. (1997). *Narrating the Organization: Dramas of Institutional Identity*. University of Chicago Press.
- Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6), 469–493.
- Dunbar, R. I. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences*, 16(4), 681–735.
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88(4), 587–604.
- Maclean, N. (1992). *Young Men and Fire*. University of Chicago Press.
- Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247.
- Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Nonaka, I., Toyama, R., & Konno, N. (2000). SECI, *ba* and leadership: A unified model of dynamic knowledge creation. *Long Range Planning*, 33(1), 5–34.
- Palermos, S. O. (2021). Distributed cognition: Bees do it and psychophysical laws. In M. D. McNeese, E.

- Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 5, p. 115). CRC Press.
- Riedl, C., Kim, Y. J., Gupta, P., Malone, T. W., & Woolley, A. W. (2021). Quantifying collective intelligence in human groups. *Proceedings of the National Academy of Sciences*, 118(21).
 - Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of Group Behavior* (pp. 185–208). Springer.
 - Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
 - Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.
 - Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688.

Chapter 9

Chapter 9 — Structure as cognitive architecture

The projection of the model in our heads

In *The Heart of Enterprise*, published in 1979, Stafford Beer breaks off his argument for one of the book's dialogue passages and gives a speaker this line: "we can only control a system to the extent that we understand it. ... all we're controlling is the projection of the model-in-our-heads."¹ It is easy to read as a remark about managers and their limits. Read it as a remark about organizations and it becomes the question of this chapter. If all an organization controls is the projection of a model, then the structure of the organization settles three things: in whose heads, and in which artifacts, that model is held; through which channels the projection is made; and how much of what comes back can be read. Roles, layers and channels are usually drawn as a chart of who reports to whom. They are also, and for this book first of all, an architecture of understanding. They decide what the organization can know.

Chapter 7 gave the unit that thinks; chapter 8 gave the conditions under which a group of such units understands more or less than its members. This chapter goes up a level. It takes its vocabulary from Beer's Viable System Model, because that model was built as an account of regulation, and a regulator, as the chapter argues, has to know. It takes from Hutchins what the model leaves out: what happens inside a channel. A companion book, *Punch and Judy*,

¹Beer (1979), p. 80.

examines the model's cybernetic inheritance in detail, and this chapter draws on its second, third, fourth, fifth and twenty-fifth chapters without restating them.²

Division of labor is division of cognition

Hutchins' formula for thinking, "the propagation of representational state across a variety of media", was introduced in chapter 1 and worked through in chapter 7.³ Apply it to an organization chart. Every box on the chart is a place where representations are held and worked on; every line is a medium across which they propagate; every level is a point at which they are re-coded. A division of labor is, whatever else it is, a division of cognition. It settles who holds which part of the picture, who sees which cues, and how far a cue has to travel before it reaches someone whose model can make something of it.

Three properties follow, and each is familiar to anyone who has worked in a large organization, though rarely named in these terms.

Spans are spans of attention. A manager with twelve direct reports and one dashboard attends to twelve summaries; the cues that do not make it into a summary do not exist at that level. Chapter 2 treated attention as the organization's first filter. Structure is where that filter is built in: the span settles how much is attenuated at each node before anything is passed up, and a span that is wide for reasons of cost is also wide as a filter.

Layers add latency. A cue that crosses three layers on the way to someone who can act on it arrives three re-codings later, and each re-coding takes time, a meeting, a report cycle. How fast a cue reaches a model that can use it is a property of the architecture, and it is the property most often blamed on people ("nobody told me") when it was

²*Punch and Judy* (manuscript, 40-stories/projects/punch-and-judy/parts/), ch. 2 "System Four Is Not a Pair of Eyes" and ch. 3 "The Law We Quote Is Not the Law Ashby Wrote"; ch. 4, 5 (the boundary) and 25 (the transducer).

³Hutchins (1995), introduction (chapter summaries; the local functional system).

drawn into the chart. The book's thread on temporality has its structural form here: the rhythm at which a level re-orientes (chapter 5) is bounded from below by the time its cues take to arrive.

Layers filter and distort. At each crossing the representation is translated into the terms of the next level, and what the next level has no distinction for is lost. Ashby's observation that "the observer and his powers of discrimination may have to be specified" before a set's variety is defined (chapter 1) becomes, in a layered organization, the observation that every layer is an observer with its own powers of discrimination, and that what the top can see is what the layers below could distinguish.⁴

None of this is measured, and the chapter does not try. Following Ashby as *Punch and Judy* reads him, the argument about variety is a priori, a way of asking whether an arrangement could work in principle, and it does not yield figures.⁵ What it yields is a way of reading a chart: as a set of decisions about where understanding can be held and how far it can travel.

System Four: the organization's outward modeling

Beer's model divides a viable organization into five systems, which he numbers. For this chapter, two of them matter, and the relation between them. System Three runs what exists: the operations as they are, today, in their present environment. System Four faces what is coming: the environment beyond the present operations, and the organization's own future in it. Every viable unit, Beer holds, needs both and needs a balance between them; chapter 1 already used this to say that understanding has to be produced at every level that faces its own environment.⁶

⁴Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

⁵*Punch and Judy*, ch. 3, note 5, on Ashby (1956), p. 208 (section 11/10).

⁶Beer, S. (1979). *The Heart of Enterprise*. Wiley; Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley. The five systems, the Three-Four

System Four is where the internal model of chapter 1 has its structural home, and where it is most often lost in teaching. The Viable System Model is commonly taught with System Four as environmental scanning: the organization's eyes, turned outward, gathering intelligence and passing it in. *Punch and Judy* calls this the “pair of eyes” reading and argues that it removes from what is taught the thing Beer put at the center.⁷ In Beer's own text System Four is the modeling and anticipating function. “System Four must contain a model — some model — of the corporation”, he wrote in *Brain of the Firm*; chapter 1 quoted the passage in full, with its consequence that a corporate plan must continuously abort.⁸ In *The Heart of Enterprise* he says it again from the other side: System Four “must contain a model of the viable system of which it is the System Four”.⁹ And foresight, in the *Brain*, is “rapid-time simulation”, which is a model run forward.¹⁰

Eyes gather; a model anticipates. The difference is the difference between an organization that receives reports about its environment and one that can say what its environment will do next and be surprised in a specific way when it does not. The first has data. The second has a projection, and a projection can be wrong, which is what makes it useful. Two cautions from *Punch and Judy* go with the point. A search of Beer's three main books finds no passage that calls the line from System Four to command a “prediction channel”, and the words “forward model” do not occur in them; the predictive-processing vocabulary that fits Beer so well is a comparison this book makes, and

homeostat and recursion are standard VSM material; the description of the homeostat follows *Punch and Judy* ch. 2 and the book's own reading.

⁷*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

⁸Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

⁹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

¹⁰Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

Beer's own claim is only what is quoted above.¹¹ And the quoted sentence has a hedge in it, "some model", which Beer meant. As the rest of the passage says, the model need not be formal, cohesive or even recognized as a model, only present.

The homeostat between Three and Four

What Beer built between Three and Four is a homeostat: two systems, each with its own model, each pulling the organization its way, held in balance so that neither wins. The balance is the design. Three has a model of the inside as it is and of what the operations can carry; Four has a model of the outside as it is becoming and of what the organization would have to be to survive there. If Four's model is allowed to run the organization, it plans for an environment the operations cannot yet serve. If Three's model runs it, the organization gets steadily better at an environment that is leaving. Where the two meet is where the organization's outward modeling meets operational reality, and where most of what this book calls orientation (chapter 5) happens at the level of the whole.¹²

Beer gives, in another place, an image of what it means for two models to meet. Writing of pupil and teacher, each of whom builds a model of the other, he says that when "the two models of each other converge, learning has occurred".¹³ The passage is about learning, and I carry it to the homeostat as my reading: the Three-Four exchange is productive when the two models converge on an account of what the organization is and can become, and it is dead when one side has stopped modeling the other. The most

¹¹*Punch and Judy*, ch. 2 "System Four Is Not a Pair of Eyes", section "Action selection: the function the diagram omits" (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

¹²Beer, S. (1979). *The Heart of Enterprise*. Wiley; Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley. The five systems, the Three-Four homeostat and recursion are standard VSM material; the description of the homeostat follows *Punch and Judy* ch. 2 and the book's own reading.

¹³Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

common way for it to die is that Four's model is never confronted with Three's at all, because the people who hold it (a strategy function, a board, an innovation lab) never sit in a forum with the people who run the operations. Chapter 8's forums were about groups; the homeostat is the same problem at the level of structure, and the commitment forum of chapter 6 is one place where it can be held.

Who carries the exchange is an open question in the tradition and in this book. Nonaka proposes "middle-up-down" management, in which middle managers link top-level vision to frontline reality. It is a candidate answer to the question of who sits in the homeostat, and chapter 10 takes it up as a question of leadership.¹⁴

Transducers, and what Hutchins adds

Every line on the chart that crosses a boundary between two systems needs, in Beer's terms, a transducer: something that converts a representation from the language of one side into the language of the other. A monthly report is a transducer. So is a manager who translates a team's complaint into a risk item, and so is the ticket system that turns a customer's sentence into a category. Beer treats transducers in terms of variety, as what a boundary lets through and what it attenuates. *Punch and Judy* discusses the transducer in its twenty-fifth chapter and the boundary in its fifth, and both come back to the choice its third chapter calls the partition problem: someone decides which states count as different, and that decision fixes what can cross.¹⁵

Hutchins supplies what happens inside a transducer. On the *Palau*, a bearing read off a gyrocompass scale

¹⁴Nonaka (1994), about pp. 29–30 (middle-up-down management, citing Nonaka 1988b). p. 30: "top management provides 'visions for direction' ... Middle management translates these visions into middle-range visions ... while top management articulates the dreams of the firm, lower managers look at the reality. The gap between these two forms of perspectives is narrowed by and through middle management."

¹⁵*Punch and Judy*, ch. 3, section "Who counts the states"; boundary in ch. 5, transducer in ch. 25.

reaches the bearing recorder as “strings of spoken digits”, is written into a log, and becomes a line on a chart (chapter 7).¹⁶ At each step the representation is re-coded for the next medium. The re-coding is the transducer at work, and it can be watched. Three things can happen to meaning on the way. It can be translated: the digits carry the bearing intact into a medium where it can be combined with other bearings. It can be lost: a digit misheard, a category on the ticket form that has no place for what the customer said. And it can be corrected: the plotter sees that the line does not cross the others where it should and asks for the bearing again. Chapter 8’s overlapping knowledge is what makes the third possible. The tissue only moves if the next person can see that the representation is wrong.

The organizational transducers that matter most are the ones nobody designed as such. A form, a template, a set of reporting categories, a KPI definition: each fixes which distinctions survive the crossing, and each was drawn by someone, for a purpose that may be years gone. Hutchins observes that the implementations of tasks are “part of a cultural process that tends to collect representations that permit tasks to be performed by means of simple cognitive processes”.¹⁷ That is what a good transducer does: it makes the next step easy. It is also why a transducer outlives the picture it was built for, and goes on making the old distinctions easy after the environment has changed. Beer’s warning that the model “must not be confused with the variety in the information that happens to be flowing” applies here with force, and his example on the same page is the case in one household: parents who see the symptoms of their son’s drug use and cannot recognize them.¹⁸ Information flows through every transducer. Whether what flows can be read depends on the model on the receiving side.

¹⁶Hutchins (1995), ch. 4, p. 226.

¹⁷Hutchins (1995), ch. 3, p. 155.

¹⁸Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

Episode — The contracts that arrived as “price”

“We have no complaint about your technicians. The other offer guarantees uptime, and they watch the units remotely. You come when we call.” The facility manager of an office-park operator said this to a Halvern account manager in the year after the expensive winter, in the call that ended a service contract covering several sites. The account manager wrote the sentence into her call notes more or less as it was said.

Then she had to close the contract in the CRM. The record could not be saved without a lost reason, and the dropdown offered four: price, competitor, site closed, other. The site was open. “Competitor”, as the sales team used the word, meant that an operator had replaced Halvern equipment with another make, and here the heat pumps stayed where they were. “Other” drew a query from sales operations. The competing contract did cost less for what it promised, so she picked price. Over the following months account managers in other regions heard versions of the same sentence from other facility managers and made the same choice.

The dropdown was the transducer at the first crossing. It had been set up years earlier, when service contracts were lost mostly to local firms that undercut Halvern on the hourly rate, and for that environment four entries had been enough. The call notes stayed in the account record. The regional sales review did not read account records. It read a table of lost contracts by reason, and in that table the facility manager’s sentence was one more count in the price row. The review’s summary for headquarters read: “price pressure in service contracts.” At this second crossing nothing was lost. A table of counts was translated accurately into a phrase the next level could work with. The loss had happened one step earlier, and the phrase carried no trace of it.

Correction would have needed someone further up who could see that the representation was wrong, and “price pressure” did not look wrong. The members of the executive strategy review knew the service business as a business of rates and margins. The phrase fitted what they knew, and nobody asked for the calls behind the counts. They answered with a discount on renewals. Renewals did not recover: operators who wanted an uptime guarantee and remote monitoring were offered the old contract for less.

None of this surprised the account managers. Each of them spoke to building operators every week and could have said what those operators were now comparing Halvern with. That account existed in the regions, in call notes and in conversations between colleagues, and no decision about the service contract depended on it.

Recursion

Beer’s structural claim that matters most for this book is recursion: a viable organization contains viable organizations, each of which has the same five systems, and each of which faces an environment of its own. Chapter 1 drew the consequence for understanding, and it can now be stated in structural terms. Every unit that faces its own environment needs its own System Four, which is to say its own model of that environment and of itself, held somewhere that the unit’s own decisions can depend on it. A picture of the whole company’s environment, held at the top, is the wrong picture for the sales region, the product team or the plant, because their environments differ from the whole’s and from each other’s, and it hides what each of them needs to see.¹⁹

¹⁹Beer, S. (1979). *The Heart of Enterprise*. Wiley; Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley. The five systems, the Three–Four homeostat and recursion are standard VSM material; the description of the homeostat follows *Punch and Judy* ch. 2 and the book’s own reading.

Two errors follow from ignoring this, and organizations commit both. The first is to hold one model at the top and to treat the units as executors of it. Their System Fours are then trivial, and the chart says so, because nothing they decide depends on any model of their own. The second is to let every unit model freely and to build no transducer between its model and the next level's, so that the organization's understanding is a set of local pictures that never meet. The homeostat between Three and Four exists at every level of recursion, and so does the question of who carries it.

Nonaka describes something similar from the side of knowledge creation. Knowledge, in his account, moves in “an upward spiral process, starting at the individual level moving up to the collective (group) level, and then to the organizational level, sometimes reaching out to the interorganizational level”.²⁰ The fit with recursion is mine, and Nonaka does not use the VSM; but the two pictures agree on the point that matters here, which is that the levels are real, that each has its own work to do, and that the crossings between them are where the work is lost.

The internal model as the core of System Four

The book has used the internal model in two senses since chapter 1, and chapter 5 settled where it uses which. In the loose sense every organization acts on some account of its situation, and the only questions are where it lives, who drew its distinctions and how it is revised. In the strong sense a regulator needs a model of what it regulates, and steering is no better than the model. This chapter is where the strong sense meets structure, and the meeting has to be done carefully, because the tradition has claimed more than it can show.

Beer claimed the most. “The regulator is as good as its contained model of the system regulated, and no more”,

²⁰Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

he wrote, and he called it a “cybernetically proven fact”.²¹ The proof he had in mind was Conant and Ashby’s 1970 paper, whose title says that every good regulator of a system must be a model of that system. The paper does not prove its title. Its authors define a good regulator as one that minimizes the entropy of the outcomes, so that a regulator which pins the system to a constant catastrophic state counts as optimal; *Punch and Judy* calls the result a proof about “a *reliable* regulator, not a desirable one”. The authors concede in the paper that not all optimal regulators are models. The proof mentions no inputs, channels or noise, and it holds only where the disturbances change “slowly with time”, which excludes the non-stationary environments System Four exists for.²²

What the book can use is smaller and holds. Francis and Wonham’s internal model principle of 1976 is the rigorous form of the claim, for control problems with stated dynamics; an organization is not obviously one, and the principle is cited here within its scope.²³ The law of requisite knowledge is a second route: having enough possible actions is not enough, the regulator must also know which one fits the disturbance.²⁴ And a reconstruction by Virgo and colleagues in 2025 replaces *being* a model with *having* one: an observer may attribute to a regulator beliefs that update consistently with the system it regulates, at the price that the attributed model may be trivial.²⁵

²¹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

²²*Punch and Judy*, ch. 3, section “The theorem that is not a theorem”, on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors’ concessions at p. 96.

²³*Punch and Judy*, ch. 3, on Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica*, 12(5), 457–465.

²⁴*Punch and Judy*, ch. 3, note 14, on Heylighen, F., & Joslyn, C. (2001). Cybernetics and second-order cybernetics. In R. A. Meyers (Ed.), *Encyclopedia of Physical Science & Technology* (3rd ed., vol. 4). Academic Press.

²⁵*Punch and Judy*, ch. 3, on Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326.

That price is what this chapter needs. Beer's "some model" and the observer's attributed model say the same thing from two sides: the model of System Four does not have to be written down, and may be "disseminated, in separate chunks of cognition around the firm".²⁶ What makes it a model rather than a figure of speech is that something depends on it. The test *Punch and Judy* proposes turns the oldest complaint about System Four, that it scans and nothing downstream moves, into a claim that can be examined: "the model attributable to this regulator is trivial."²⁷ A System Four whose attributed model is trivial is one on which nothing decided depends. It produces reports; it changes no decision; the organization would act the same without it. That is the testable failure of the function, and it is visible in the minutes: find the decisions of the last year, and ask which of them would have gone differently had the outward model said something else.

Requisite variety belongs here too, and belongs here as Ashby wrote it. "Only variety in R can force down the variety due to D; variety can destroy variety."²⁸ The sentence usually quoted, "only variety can absorb variety", is Beer's compression, from *Brain of the Firm*, and it is the better sentence and the more dangerous one, because it makes a design instruction out of a law and travels without the law's conditions.²⁹ The law is information theory's channel capacity applied to regulation: "R's capacity as a regulator cannot exceed R's capacity as a channel of communication".³⁰ It is a priori and licenses an argument about whether an arrangement could work, never a number; the one careful attempt to make it measurable

²⁶Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

²⁷*Punch and Judy*, ch. 3, section "From a complaint into a claim", on Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025), A "good regulator theorem" for embodied agents, arXiv:2508.06326. (P&J).

²⁸Ashby (1956), p. 207 (section 11/7), pagination as in *Punch and Judy*, ch. 3. "Only variety can absorb variety" is Beer's compression (*Brain of the Firm*, 2nd ed., p. 308).

²⁹*Punch and Judy*, ch. 3, "Two sentences and one verb" and note 2, on Beer, *Brain of the Firm* (2nd ed., 1981), p. 308.

³⁰*Punch and Judy*, ch. 3, note 4, on Ashby (1956), pp. 209–211 (section 11/11).

ended with a judgment of adequacy rather than a count of states.³¹ For this chapter that means the question to ask of a System Four is whether the model it holds can make the distinctions its environment will present it with, and the answer is an argument.

Beer takes the modeling one level further, and the book follows him. System Five, the unit's policy and identity, "must necessarily contain models of System Three and Four".³² Beyond modeling its environment and itself, the organization models its own modeling: it holds an account of what its operations can do and of what its outward function believes, and it arbitrates between them from that account. The good-regulator reading adds a limit that Five has to respect. Model quality says nothing about whether the state being defended is worth defending; that question falls to identity and purpose, and no amount of modeling settles it.³³

Structure itself can be read as a model

If the model may be disseminated in chunks and need not be recognized as a model, then it may be held in the structure. Raúl Espejo and Alfonso Reyes make the structural reading explicit: an organization's structure, they write, "should map" the complexity it chooses to respond to.³⁴ The units an organization has, the environments it assigns them, the channels it builds between them: these are its claim about what its environment consists of and which of its distinctions matter. An organization with a regional structure has modeled its world as one in which

³¹Punch and Judy, ch. 3, section "What the law becomes when someone has to measure it", on Pfiffner (2017), p. 165; the a-priori character of the law, Ashby (1956), p. 208 (section 11/10). (P&J)

³²Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

³³Punch and Judy, ch. 3, section "The theorem that is not a theorem", on Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97: optimality criterion at p. 91, the authors' concessions at p. 96.

³⁴Espejo, R., & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Springer, p. 95; see also p. 138.

geography is the distinction that counts; one that reorganizes around customer segments has revised the model, whatever the announcement says. (Espejo and Reyes invoke the Conant–Ashby title as a theorem on the way; the structural reading does not need it, and the book takes the reading without the theorem.)

Hutchins gives the structural model its history. The artifacts and arrangements of a working system accumulate the solutions of earlier generations: “Humans create their cognitive powers by creating the environments in which they exercise those powers.”³⁵ A structure is such an environment. It holds the distinctions that once fitted the environment and made the work easy, and it goes on holding them. That is capacity, and it is inertia. The structural model is the slowest of the organization’s models to revise, because revising it means moving people, and because nobody thinks of it as a model. Chapter 11 returns to this as the part of unlearning that organizations skip.

There is a corollary that Hutchins states for tools and that holds for structure. “None of the component cognitive abilities has been amplified by the use of any of the tools. Rather, each tool presents the task to the user as a different sort of cognitive problem requiring a different set of cognitive abilities”.³⁶ A restructuring does not amplify the organization’s understanding. It changes the task of understanding: which cues are near whom, which crossings a representation has to survive, whose model a decision depends on. That is why structural changes that alter channels and spans should change the quality of understanding, and why, if they were found not to, the argument of this chapter would be in trouble.

What the diagram omits

The Viable System Model has a gap that this book met from the side of chapter 6, and *Punch and Judy* names

³⁵Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

³⁶Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

it: action selection, “the function the diagram omits”.³⁷ Among the actions available to a unit at any moment most are incompatible with each other, and “something has to decide which one is expressed”. The diagram shows where models are held and where variety is exchanged; it does not show how a field of possible actions is “held back, compared, and selectively released”. *Punch and Judy* draws the comparison with the vertebrate case, where selection works by releasing one candidate from inhibition; the book uses it as an analogy only, as chapter 6 did.³⁸

For the architecture, the gap matters in one specific way. If action selection is a function, it has to sit somewhere, and where it sits is a structural decision the diagram does not force. Beer’s own answer, given elsewhere in the *Brain*, is the redundancy of potential command: “the command centre changes from time to time. Its location is a function of the information available”, and in a rigid pyramid “there can be no redundancy of potential command”.³⁹ Read with this chapter, that means selection should be released where the relevant model is held, at whatever level of recursion that is, and that a structure which fixes the point of selection at the top has separated the choice from the model that could make it. Hutchins shows the extreme case, in which the selection was made by no one: the *Palau*’s new division of labor was “evolutionary in the sense that the system-level change that resulted was never represented”.⁴⁰ A structure can select without a selector. Whether it should, and who answers for what it selects, is chapter 10’s question.

³⁷*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

³⁸*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02 system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

³⁹Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

⁴⁰Hutchins (1995), ch. 8, p. 349.

The inter-organizational level

The environment that System Four models contains other organizations, each with a System Four of its own, each modeling an environment that includes us. Chapter 1 said this in passing; the structural reading gives it an edge. A model of the environment that treats customers, competitors and regulators as conditions rather than as models will be surprised in a way it cannot explain, because the environment will have anticipated it. The transducers at the organization's outer boundary, the contracts, the standards, the account each side keeps of the other, are two-sided in a way the internal ones are not: what crosses them was drawn by two observers, and the partition problem has two authors.

And there are units above the organization that face environments of their own. Networks, ecosystems, supply chains and industries carry the recursion question whether or not anyone has drawn them as viable systems: is there a System Four at that level, whose model is it, and does anything decided at that level depend on it? Nonaka's spiral "sometimes reaching out to the interorganizational level" describes knowledge crossing that boundary.⁴¹ The book has no more to say about the level than that the question exists, and that most of the organizations in a network answer it, by default, with nobody.

What would count against it, and what comes next

The chapter's claims are structural, and they can be tested structurally. If changes that alter channels, spans and layers left the quality of understanding unchanged, the reading of structure as cognitive architecture would be wrong. If a unit whose System Four holds a rich, current model steered no better than one whose attributed model is trivial, the strong sense of the internal model would be wrong at this level, whatever its warrant. Where to look: the diagnostic cases of the VSM literature, Espejo and

⁴¹Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

Reyes among them; and the accident analyses of chapter 1, in which cues crossed layers and were lost, which are the natural experiments nobody wanted.⁴²

The projection of the model in our heads, then, is made through a structure, and the structure settles whose heads. Spans and layers settle which cues reach a model at all; transducers settle which distinctions survive the crossing; recursion settles at which levels a model is held; and the whole arrangement is itself a model, drawn by someone. Structure decides who draws the distinctions a unit can see. Chapter 10 asks who holds that power, and what a leader's job in understanding is once the picture no longer belongs to the leader.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Beer, S. (1985). *Diagnosing the System for Organizations*. Wiley.
- Cole, M., & Griffin, P. (1980). Cited in Hutchins (1995), ch. 3.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
- Espejo, R., & Reyes, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*. Springer.
- Francis, B. A., & Wonham, W. M. (1976). The internal model principle of control theory. *Automatica*, 12(5), 457–465.
- Heylighen, F., & Joslyn, C. (2001). Cybernetics and second-order cybernetics. In R. A. Meyers (Ed.), *Encyclopedia of Physical Science & Technology* (3rd ed., vol. 4). Academic Press.

⁴²Espejo & Reyes (2011), diagnostic cases (the book). Vaughan, D. (1996). *The Challenger Launch Decision*. University of Chicago Press. Snook, S. A. (2000). *Friendly Fire*. Princeton University Press.

- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Pfiffner (2017).
- *Punch and Judy*.
- Snook, S. A. (2000). *Friendly Fire*. Princeton University Press.
- Vaughan, D. (1996). *The Challenger Launch Decision*. University of Chicago Press.
- Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326.

Chapter 10

Chapter 10 — Whose account counts: leadership and sensegiving

Who really decides to buy

On page 232 of the second edition of *Brain of the Firm*, published in 1981, Stafford Beer stops his argument to ask a question that most organization charts answer wrongly: “Who, for example, really decides to buy”. His answer is not the director whose signature is on the order. “It is the little group of employees, possibly very junior, which understands the need”.¹ The signature ratifies an account of the situation that was built somewhere else, by people who saw the need before anyone with authority did. Beer took the idea from Warren McCulloch, who had called it the redundancy of potential command. In a system built to survive, “the command centre changes from time to time. Its location is a function of the information available”.² Command follows information. Beer draws the structural consequence on the same page: in a rigid pyramid “there can be no redundancy of potential command”.³

I begin with the junior group because the question of this chapter is easy to state and hard to keep in view. Who shapes the account an organization acts on? Chapters 7 to 9 said where understanding is held, in a system of people, artifacts and structure; when a group holds more of it than its members; and how spans, layers and transducers

¹Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

²Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

³Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

filter it on its way. None of that said whose version wins when two accounts of the same situation are in the room. Beer's junior group understood the need. Whether the organization acted on their understanding, and in whose words, is a different matter, and it is a matter of power.

Edwin Hutchins, whose navigation team carried Part III, is of no help here, and he does not claim to be. Rank and military identity appear in his book; who decides which account counts does not.⁴ His cases are well-structured work in which the situation is a matter of computation, and the ship's position is not a topic on which the captain and the plotter can hold rival views for long. The literature that treats accounts as contested comes from Karl Weick's side, and the chapter follows it there, and then back to Beer.

Sensegiving

In 1991 Dennis Gioia and Kumar Chittipeddi published a study of the beginning of a strategic change at a large public university, followed from inside during a new president's first year in office.⁵ Their lasting contribution was a word. Beside *sensemaking*, which by then carried Weick's meaning, they set *sensegiving*: the attempt to influence how others make sense, to move the meaning others construct toward a preferred account of the organization and its situation. In their reading the two alternate. The president and his team first had to make sense of the university for themselves; then they had to give that sense to others, whose reactions changed what the leaders had to make sense of next.

Two things follow that the earlier chapters left implicit. First, sensemaking in organizations is political from the

⁴Hutchins, E. (1995). *Cognition in the Wild*. MIT Press. Rank and military identity in ch. 1; power and politics are not a theme.

⁵Gioia, D. A., & Chittipeddi, K. (1991). Sensemaking and sensegiving in strategic change initiation. *Strategic Management Journal*, 12(6), 433–448. p. 433: “an ethnographic study of the initiation of a strategic change effort in a large, public university ... tracking the first year of the change”; “the emergent concepts of ‘sensemaking’ and ‘sensegiving’”; p. 435: “The research project began with the arrival of a new President”.

start. The account an organization ends up acting on is rarely the one that best fits the cues. It is the one that was given most effectively, by people with the standing and the channels to give it. Leaders give sense as much as they make it. Second, giving sense is part of a leader's work whether or not the leader intends it. A leader who says nothing about a surprise has given the sense that it does not matter. Chapter 4 described sensemaking as the social construction of a plausible account. The sensegiving literature adds that the construction has a foreman, and the foreman is not always the person with the best view of the site.

The medium in which sense is given is the one chapters 4 and 8 named: narrative. An account travels between people as a story, and stories compete on terms that have little to do with fit. The story that is easiest to retell, that has a villain, that explains the last three quarters in one sentence, wins against the story that is true in four. A leader's advantage in this competition is structural rather than rhetorical. The leader's story is told first, from the front of the room, and it is repeated by the people whose standing depends on repeating it; the junior group's story is told in the corridor, once, to whoever will listen. Sensegiving from above needs no talent to succeed. It needs only the absence of a forum in which the other story could be told to the same audience.

Sally Maitlis followed the point into detail. In a study of three British symphony orchestras she distinguished four forms that the social process of sensemaking takes, according to how much the leaders and how much the other stakeholders engage in sensegiving.⁶ Where both do, the process is *guided*: it produces a single account, and a rich

⁶Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21-49. p. 21: "four distinct forms: guided, fragmented, restricted, and minimal. These forms result from the degree to which leaders and stakeholders engage in 'sensegiving'"; p. 24: three British symphony orchestras; p. 32 (Fig.): guided → "Unitary, rich account". The two dimensions are control and animation, so "how much leaders and stakeholders control them" was a simplification, corrected in the text.

one, because it has been argued over. Where stakeholders give sense and leaders do not, it is *fragmented*: many accounts, each thin, and no organizational one. Where leaders give sense and stakeholders do not, it is *restricted*: one account, narrow, never tested against what the front line knows. Where neither does, it is *minimal*: a nominal account that decides nothing. The labels are hers; the glosses are my recollection. The shape of the finding matters more than the labels. The form of the process decides which account the organization ends up with, and the form is something leaders choose, mostly by what they do not do.

Set the four forms beside Beer's page. The restricted form is the rigid pyramid: leaders hold the interpretation, the group that understands the need is not asked, and there is no redundancy of potential command. The fragmented form is its mirror: many junior groups understand many needs and nobody joins the accounts up. The guided form is the one in which the command center can move to where the information is and still return with an account the whole organization can act on. Maitlis's stakeholders and Beer's junior employees are the same people seen from two disciplines, and the two disciplines agree on what happens to them in a pyramid.

Who draws the distinctions

There is a level below sensegiving at which power enters the account, and it is easy to miss because nothing is said there. Chapter 1 quoted Ross Ashby: "a set's variety is not an intrinsic property of the set: the observer and his powers of discrimination may have to be specified".⁷ Which states count as different is decided by an observer. *Punch and Judy* calls this the partition problem. Somebody partitions the world into the distinctions a model can make, and from that moment the model sees only the differences it was built to see; the observer, and with the

⁷Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

observer a politics, has entered the model before any data arrives.⁸ Beer's example, which chapter 2 used, is the parents who cannot read the signs of their son's drug use. The information is flowing past them, and their model has no distinction for it.⁹

In chapters 2 and 5 the partition problem explained why more data does not help. Here it returns as a question of power. Whoever draws the distinctions decides what can count as a signal at all, and that person is rarely appointed to the job. A strategy team defines the market segments; from then on a customer who fits none of them is noise. A risk function defines the risk categories; a risk that belongs to none of them is not a risk. A dashboard designer chooses the axes. Each of these is an act of sensegiving that precedes any argument about what the numbers mean, and it is the most durable kind, because it is built into the instruments rather than said in a meeting. Gioia and Chittipeddi's president gave sense in speeches and documents, which people could answer. The partition is given in the form fields, the taxonomy and the reporting template, and it keeps giving sense long after the people who drew it have left the building.

Chapter 9 said that a transducer re-codes a representation for the next medium and that meaning is translated, lost or corrected on the way. The partition is what a transducer enforces. A frontline observation that has no category in the reporting template does not cross the boundary as an anomaly; it crosses as "other", or it does not cross at all. The collective form of the partition problem is therefore a question about who may contest a category. In most organizations nobody may, because the categories are not experienced as anyone's decision. They are experienced as the system.

The consequence for leadership is a question that most leaders have never been asked: who drew the distinctions your organization sees with, and when? If the answer is a consultant in 2014, or the vendor of the reporting system,

⁸*Punch and Judy*, ch. 3, section "Who counts the states"; the boundary in ch. 5, the transducer in ch. 25.

⁹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 352.

the organization's account of its situation has an author its leaders did not choose. Redrawing distinctions is possible. It is also one of the few moves that only leadership can authorize, because the partition sits in shared artifacts that no single unit owns, and a unit that redraws its own will find that its numbers no longer add up with anyone else's. That makes the partition the first item in the job description that follows.

The leader's job: enabling orientation

If the account is distributed (chapter 7), if the quality of a group's account depends on how its communication is organized (chapter 8), if structure filters what reaches whom (chapter 9), and if the account is contested (this chapter), then the picture of the leader as the person who holds the interpretation is wrong on all four counts. The leader cannot hold it. The leader's attempt to hold it degrades it, by closing the channels through which it would have been corrected. And the leader's version is one of several, with the advantage of standing and the disadvantage of distance. Beer put the limit of the single interpreter in a sentence from *The Heart of Enterprise*: "we can only control a system to the extent that we understand it. ... all we're controlling is the projection of the model-in-our-heads."¹⁰ A leader who holds the interpretation is steering by the projection of one head's model, and the head is the one furthest from the customer, the ticket and the workaround.

What remains is a job that is smaller in one way and larger in another: to enable orientation, in the sense chapter 5 gave the word. The leader is responsible for the conditions under which the organization's model is kept alive, and not for the model's contents. Four conditions follow from the earlier chapters, and I take them in turn.

Identity and purpose. The internal-model argument of chapter 5 has a limit, and *Punch and Judy* states it: how good a regulator's model is says nothing about whether the state it defends is worth defending. That question belongs

¹⁰Beer (1979), p. 80.

to identity and purpose, and no amount of modeling answers it.¹¹ Weick makes identity the first of the properties of sensemaking: who we take ourselves to be decides what we can see and what counts as a threat.¹² Nonaka names intention first among the conditions under which people create knowledge at all.¹³ In chapter 9, Beer's System Five was the place in the structure that holds the models of Three and Four and arbitrates between them. In practice this is what a leader supplies before any observing starts: what this organization is, what it is for, and what would count as winning. In my own consulting practice I have called this setting the game before the observing, and the phrase is mine rather than any of the theories'.¹⁴ Set the game and the partition has a purpose to be drawn for. Leave it unset and every unit draws its own, which is Maitlis's fragmented form under another name.

Forums in which accounts can compete. Chapter 8 argued that a group understands more than its members only under conditions, and that one of them is a forum that separates the competition between frames from the commitment to one of them. The leader's instrument is the forum. Guided sensemaking, in Maitlis's sense, has both sides giving sense in a place designed for it, with the leader accountable for the outcome and for the place. That includes the discipline of keeping the forum open after the leader has formed a view, which is the moment at which most forums close, and the moment at which the restricted form begins.

¹¹*Punch and Judy*, ch. 3, on the good-regulator reading: model quality says nothing about whether the state defended is worth defending; that question falls to identity and purpose.

¹²Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, ch. 2, identity construction as the first of seven properties.

¹³Nonaka (1994), p. 17 (individuals create knowledge; the organization amplifies it; intention, autonomy and fluctuation; Weick 1979 on enactment), p. 18 (breakdowns after Winograd and Flores). The gloss of "breakdown" as the moment a routine stops working and its assumptions become visible is the author's; Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

¹⁴The formulation "the game before the observing" is the author's, from his own consulting material on the cycle described in Appendix A (the GRADO note "The Game"). Standalone readers: Appendix A.

Protection for the people who report weak signals.

Chapter 2 said that weak signals are filtered by the model, by channels and by silence. The first two are design. The third is a question of what happens to the person who speaks. Morrison and Milliken called the result organizational silence: people who see a problem and calculate that saying so is not worth what it will cost them.¹⁵ Beer's junior group understood the need. Whether they said it, and to whom, depended on what had happened the last time someone junior said something. Protection here is concrete and it is the leader's alone: who was thanked, who was contradicted in public, what became of the last messenger. Edmondson's psychological safety, in chapter 8, is the team-level name for the same condition, and it is set from above more often than team-level research suggests.

Decision rights where the model sits. Chapter 6 argued that for experienced people understanding already contains the action, and that an organization therefore has to place decision rights where the relevant model is held. Beer's redundancy of potential command is the structural form of the same rule, and it is where sensegiving and action selection meet. If the junior group understands the need and the decision sits three levels up, the organization has two options: move the decision to the group, or move an account of the need up three levels through three transducers, losing something at each. Beer's answer is the first, and the rigid pyramid is the design that forbids it. Hutchins supplies the extreme case. On the *Palau* the new division of labor was, in his words, "evolutionary in the sense that the system-level change that resulted was never represented". Nobody decided it, and it worked.¹⁶ That is not an argument for never deciding. It is a reminder that the account and the choice can move to where the information is without anyone giving permission, and that

¹⁵Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725. p. 706 (organizational silence defined), p. 708 (a climate of silence: futile and/or dangerous).

¹⁶Hutchins (1995), ch. 8, p. 349.

a structure which cannot let them move will be steered by a model that is out of date by the height of the pyramid.

A fifth condition runs across the four, and it is the temporality thread of the book. Chapter 5 said that how often a model is revisited is a design variable. Who sets it is a leadership question. The calendar of forums, the interval between strategy reviews, the point in a project at which someone is allowed to say that the premise was wrong: each of these is set from above, usually without anyone noticing that a rhythm has been chosen. An organization whose leaders revisit the account once a year has told the junior group how long a need will wait before anyone with authority hears of it. Rhythm is sensegiving by the calendar.

Four conditions and a rhythm, and none of them is the interpretation. A leader who supplies all of them and no interpretation of her own has done the job. A leader who supplies only the interpretation has done the part of it that the organization could have done without her, and has usually done it in the restricted form.

The middle

Between the leader who sets the game and the junior group that understands the need sits a layer that both the sensemaking literature and the knowledge-creation literature treat as the hinge. Ikujiro Nonaka, whose theory carries Part IV, proposed what he called middle-up-down management. Against top-down, in which the top's vision is handed down for execution, and bottom-up, in which frontline knowledge is expected to rise on its own, he placed middle managers as the people who translate in both directions: a vision into concepts the front line can work with, and frontline knowledge into something the top can steer by.¹⁷ I cite this from the later pages of his 1994 article (p. 30).

¹⁷Nonaka (1994), about pp. 29–30 (middle-up-down management, citing Nonaka 1988b). p. 30: “top management provides ‘visions for direction’ ... Middle management translates these visions into middle-range visions ...”. Not a model summary any longer.

What I can verify is older. In 1986 Hirotaka Takeuchi and Nonaka described the “rugby” approach to product development in Japanese firms, in the article from which Scrum later took its name, and among its six characteristics they listed self-organizing project teams and what they called subtle control, side by side.¹⁸ My gloss of the pair: management that sets the direction and the boundaries and leaves the team to find the route. The teams in that article are the junior group with room to act. Subtle control is the middle’s job, to keep the game set from above and the account coming from below in one frame, without doing either side’s work for it.

In the vocabulary of chapter 9, the middle is where the model of the environment and the model of the organization have to be held together, and where the transducers on both sides are operated by the same people. In the vocabulary of this chapter, the middle is where restricted sensemaking is manufactured. A middle manager who reports upward what the top expects and downward what the top decided has closed both channels while appearing to operate them, and no forum at the top will find out, because the forum is fed by the same manager. The redundancy of potential command depends on the middle letting command move through it, in both directions, and a middle that has learned that its standing depends on being the channel will not.

Identity under threat

Chapter 1 told the Mann Gulch story in full, and I return to it only for the point that belongs here. When the fire turned on the crew, what collapsed was not only the picture of the fire. The role structure went with it. Weick’s term for the moment is a cosmology episode, and he

¹⁸Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116, p. 2 (the six characteristics). That Scrum took its name from the article’s rugby metaphor rests on the article itself — its section heading is “Moving the Scrum Downfield” — and on the six characteristics listed above.

adds that a crew put together for the job is “susceptible to sudden losses of meaning”, because it has no shared structure beneath the roles in which a new picture could be built.¹⁹ The foreman’s escape fire was an action whose meaning the crew could not construct in time. Read with this chapter, it was sensegiving without a forum, delivered to people who had just stopped being the organization it was addressed to.

The general point is about identity, and it is the chapter’s thread. Sensegiving works on people who know who they are, in an organization they recognize. When a surprise threatens identity, the first thing that goes is the capacity to receive an account, however good. This is why leaders find that the interpretation they give in a crisis does not land: they are giving sense to a structure that no longer exists to hear it. Rebuilding the role structure, the sense of who is who and who does what, comes before any new picture, and it is leadership work in a sense that no forum design covers. Weick’s soldiers with the wrong map, in chapter 5, found their way because they had a purpose and an image of where they were going; the map came second. The leader’s contribution when the frame breaks is the purpose and the roles. The map can wait, and usually has to.

Most of that work is done before the surprise, and the *Palau* shows what it consists of. The navigation team’s role structure survived the loss of the gyrocompass because the roles overlapped. Hutchins found that overlapping knowledge between roles is “an important source of the robustness of such systems in the face of error and interruption”, and that because of it the team could “reconfigure dynamically”; in his image, “the individuals are a sort of flexible tissue that moves to ensure the propagation of task-relevant representational state”.²⁰ A pyramid treats overlap as waste. Every role is defined so that no two people know the same thing, and the redundancy that would

¹⁹Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652. p. 633.

²⁰Hutchins (1995), ch. 4, pp. 221 and 227.

let the structure survive the loss of a part is removed as an efficiency. The leader who wants an organization that can rebuild its roles under threat has to have paid for the overlap in advance, when it looked like slack, and to have defended it against the people whose job was to remove it.

The emotional side belongs here too, and the book has said less about it than it deserves. The collapse of a frame is felt before it is understood, and fear narrows what can be built; the evidence for both is in the crisis-sensemaking literature.²¹ A leader who gives a calm account into fear is not lying. The account will be received by a narrower set of people than the organization chart suggests, and the leader who does not know that will conclude that the account was wrong, when it was only early.

The dark side

Sensegiving has a darker side, and it works by force as much as by argument. An organization can be made to feel that its old picture no longer holds so that a preferred one can replace it, or the cues that would unsettle a preferred picture can simply be kept back. Both are ordinary. A reorganization announced as the response to a crisis that nobody felt as one until it was announced works the first way. A dashboard that shows the metrics on which the current strategy looks good, and omits the ones on which it does not, works the second, and it is also the partition problem in its most deliberate form.

A chapter on whose account counts has to say how accounts are made to count by force, and this is the place. Every one of the four conditions above has a dark twin:

²¹Maitlis, S., & Sonenshein, S. (2010). Sensemaking in crisis and change: Inspiration and insights from Weick (1988). *Journal of Management Studies*, 47(3), 551–580. Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247. Maitlis & Sonenshein 2010 (p. 567: intense negative emotion impedes sensemaking; p. 570: fear narrows); Maitlis, Vogus & Lawrence 2013 (p. 222: emotion fuels sensemaking; p. 227: intense emotion forestalls it); the paper's subtitle cites Weick (1988), not 1993.

identity set as an exclusion, forums held to ratify, messengers punished, and decision rights kept precisely where the model is not. Maitlis's restricted form is what these produce when they are stable. What they cannot produce is the one thing the organization needs from its account, which is that it fit the situation better than the alternatives. A restricted account is unitary because it was never tested, and the test is postponed rather than avoided: it arrives later, from the environment, in the form of a surprise that the account had no distinction for.

Episode — The Product Quality Board

"Installation errors rising. Measure: partner training." That was the bullet in the monthly quality report at Halvern, a manufacturer of heat-pump systems for commercial buildings, in the autumn after its C4 controller generation had gone into the field. Henrik, the Head of Quality, presented it to the Product Quality Board. His department had written it from root-cause codes it assigned itself, five of them, mapped one-to-one onto the five cause categories of the service ticket: product defect, installation error, operating error, wear, other. The template of the report was his too, with one bullet per category, a trend arrow and a measure.

The Board met monthly. Dr. Reuter, the COO, chaired it, and the C4 was his product. Quality presented; Service attended and had no slot on the agenda. Since the spring, technicians in region South had been saying that some C4 units kept restarting once the building-management system was connected, and building operators had said the same on the phone to the support desk. Because an installer had made that connection, the tickets were filed under installation error. The sentence itself went into the free-text field, which no report carried, and into the technicians' messenger group, which was not a channel the Board read. Milan, a senior specialist on the second line, had seen the pattern in a handful of escalations by early summer

and had not pressed it. Nothing was withheld from the Board. Henrik showed what his codes contained, and the Board approved the measure on the page. The process was of the restricted kind: one account, given from the front of the room, ratified by the forum that owned the remedy.

The partner training program started in November and cost about €180,000. The executive board's quarterly slide attributed a 9% rise in service cost to installation quality at partners. Over the winter 96 compressors were replaced, against about 25 in a normal season, and two large installer partners, who had been told the fault was theirs and sent to training, moved part of their business elsewhere.

In March Katrin, the Head of Service, asked for an agenda slot. She did not bring a rival bullet. She brought Milan, Tobias, a technician from region South, and three cases traced unit by unit: when the building system was connected, when the short-cycling began, when the compressor failed. In each, the commissioning record was clean and the cycling started in the week of the connection. Tobias said what he had seen on site. Milan described the two parameter changes he had worked out in February and what they implied about the cause. Henrik asked which code the three cases belonged under, and nobody in the room could name one. Dr. Reuter asked whether the changes held on units already cycling; Milan said they had so far. By the end of the meeting the Board's account was an interaction fault between C4 firmware and the building system's setpoint overrides, and the question on the table was a service bulletin instead of more training.

The group that understood the need had been on the road in region South and on the second line for almost a year. Its account reached the people who could act on it when someone with a seat in the forum put it there, with its authors, in front of the same audience

that had heard the first one. The five codes were still five when the meeting closed.

What would count against it

The chapter claims that organizations in which leaders enable distributed orientation, under the four conditions above, adapt better than organizations in which leaders hold the interpretation centrally. Against it would count: organizations in which leaders hold the interpretation centrally adapting as well as those that enable distributed orientation, in comparable settings; guided sensemaking producing accounts no more robust under surprise than the restricted form; protection for the people who report weak signals changing nothing in what surfaces. Where to look: Maitlis's orchestra studies and the strategic-change case research that followed Gioia and Chittipeddi; the organizational-silence studies named in chapter 2; and, for the redundancy of potential command, the evidence on decentralized decision rights in settings where the environment moves faster than the reporting chain.

Handover: an account committed to

Return to page 232. Beer's junior group understood the need before anyone with authority did. If the organization was built well, the command center moved to them, the account they held became the account the organization acted on, and a decision was committed, with a record, in the sense chapter 6 gave the word. That is the end of Part III: an account that has been distributed, organized, structured, contested and committed to.

It is also the beginning of Part IV, and the reason the two parts join is the *Palau*. What the junior group understood, the organization acted on. What the organization learned from acting on it is a different question, and Hutchins' team shows what happens when it is not asked:

the solution “was not saved in the system”.²² Sensegiving decides which account the organization acts on. Learning decides what of the outcome the organization keeps, and chapter 11 begins there.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Gioia, D. A., & Chittipeddi, K. (1991). Sensemaking and sensegiving in strategic change initiation. *Strategic Management Journal*, 12(6), 433–448.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21–49.
- Maitlis, S., & Sonenshein, S. (2010). Sensemaking in crisis and change: Inspiration and insights from Weick (1988). *Journal of Management Studies*, 47(3), 551–580.
- Maitlis, S., Vogus, T. J., & Lawrence, T. B. (2013). Sensemaking and emotion in organizations. *Organizational Psychology Review*, 3(3), 222–247.
- Morrison, E. W., & Milliken, F. J. (2000). Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25(4), 706–725.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- *Punch and Judy*.
- Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986.

²²Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

- Weick, K. E. (1993). The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 38(4), 628–652.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.

Chapter 11

Chapter 11 — Learning loops: how organizations create knowledge, and lose it

A warning inside a success story

In the January–February 1986 issue of the *Harvard Business Review*, Hirotaka Takeuchi and Ikujiro Nonaka described how a number of Japanese companies were developing new products: in overlapping phases, by teams that organized themselves, in a manner the authors compared to rugby. The article is the one from which Scrum later took its name.¹ On the ninth page of the reprint, after eight pages on how those teams worked, the authors turned on their own material. “Institutionalization, when carried too far, can create its own danger. Passing down words of wisdom from the past or establishing standard practices based on success stories works well when the external environment is stable.”² The remedy they named was unlearning: the deliberate discarding of lessons that had once been true.

Two things in that article carry this chapter. The first is how the companies moved what they had learned from

¹Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116, p. 2 (the six characteristics). That Scrum took its name from the article’s rugby metaphor rests on the article itself — its section heading is “Moving the Scrum Downfield” — and on the six characteristics listed above.

²Takeuchi & Nonaka (1986), reprint p. 9 (institutionalization and unlearning; the limits of the sample).

one project to the next. They did it through people. Learning was transferred “through ‘osmosis’”, by assigning the people who had carried one project to the next one.³ The second is the warning: what an organization learns settles into procedures and artifacts, and there it keeps working after it has stopped being true. Between the two lies the question of the chapter. An organization acts on an account of its situation, shaped and committed to in the ways Part III described; the action produces an outcome; what does the organization learn from it, where does that learning go, and why is so much of it lost?

The chapter’s claim is that understanding is the precondition and the catalyst of organizational learning, and that the creation of knowledge is what follows it. The theory it uses to say how is Nonaka’s, set out in a 1994 article and developed with Takeuchi in *The Knowledge-Creating Company* the year after.⁴ The counterpoint is the *Palau*, and the chapter returns to the osmosis of 1986 at its end.

Understanding first, then knowledge

Nonaka’s epistemology is on the second page of the 1994 article, and it is stricter than anything in Part II. Knowledge is “justified true belief”, and the emphasis falls on “a dynamic human process of justifying personal beliefs”.⁵ Compare that with Weick. Sensemaking, in chapter 4, is driven by plausibility rather than accuracy: an

³Takeuchi & Nonaka (1986), reprint p. 8 (subtle control; transfer of learning through “osmosis”).

⁴Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 15. <https://doi.org/10.1287/orsc.5.1.14> (public PDF, pp. 14–21). Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.

⁵Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 15. <https://doi.org/10.1287/orsc.5.1.14> (public PDF, pp. 14–21). Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.

account good enough to act on, tested by the acting. Nonaka's knowledge is what has survived the test. The two standards describe the same account at two moments. Plausibility is enough to act, and justification, in Nonaka and Takeuchi's sense, is the organization's later judgment of whether the account is worth keeping, taken against the intention it has set rather than by the situation alone. Sensemaking gives the organization a hypothesis; knowledge creation is what happens when the hypothesis has been tried and someone takes the trouble to say what held.⁶

That is why understanding comes first here as well. Nonaka himself starts from the individual who is trying to "make sense of their environment", citing Weick on enactment, and names three conditions under which that individual creates knowledge: "'intention,' and 'autonomy,' and a certain level of environmental 'fluctuation'". Fluctuation he relates to breakdowns in Winograd and Flores's sense, the moment a routine stops working and the assumptions built into it become visible.⁷ That moment is Weick's interruption and Endsley's failed projection; it is where sensemaking starts in chapter 4. The same event that triggers sensemaking triggers the creation of knowledge, and an organization that suppresses surprise, in the ways chapter 2 described, has removed the occasion for both.

The order matters in practice because organizations often try to run the second without the first. A lessons-learned exercise after a project asks what was learned before anyone has made sense of what happened, and the result is a list of what the template asked for. A knowledge

⁶The reading of plausibility and justification as the same account at two moments is the author's; neither Weick's nor Nonaka's. (author's reading)

⁷Nonaka (1994), p. 17 (individuals create knowledge; the organization amplifies it; intention, autonomy and fluctuation; Weick 1979 on enactment), p. 18 (breakdowns after Winograd and Flores). The gloss of "breakdown" as the moment a routine stops working and its assumptions become visible is the author's; Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

base fills with articles about situations nobody has understood. Nonaka's word for what those articles contain is explicit knowledge, and it is one of two kinds.

Two kinds of knowledge, four movements

Chapter 1 gave the distinction. Following Michael Polanyi, Nonaka separates explicit knowledge, “transmittable in formal, systematic language”, from tacit knowledge, which is “hard to formalize and communicate” and “deeply rooted in action, commitment, and involvement in a specific context”.⁸ Tacit knowledge has cognitive elements, which are the mental models of chapter 5, and technical ones: “concrete know-how, crafts, and skills”.⁹ The service agent who hears a customer's language change, the engineer who knows which workaround will hold, the finance manager who feels the budget pressure before the report shows it: each holds knowledge of the second kind, and each would have trouble writing it down.

Nonaka's theory is a theory of movement between the two kinds. Knowledge is created, he argues, when it moves from tacit to explicit form and back, and when it moves from one person to a group and to the organization. The four movements are the ones every summary of the theory lists, and I give them in his words where the words are verified.¹⁰

Socialization creates “tacit knowledge through shared experience”. It is how apprentices learn from masters, “by observation, imitation, and practice”, and it needs no language. It needs presence. In organizations it happens in the shift handover, the joint customer visit, the pairing of a new engineer with an old one. It is the mode that documentation cannot replace, and most organizations do

⁸Nonaka (1994), p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4. Quoted in chapter 1.

⁹Nonaka (1994), p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4. Quoted in chapter 1.

¹⁰Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

not know they depend on it until the person who carried it leaves.

Externalization moves tacit knowledge into explicit form. It is the hardest of the four and the one on which Nonaka spends most of his attention. Its engine, in his account, is metaphor: “‘Metaphor’ plays an important role in the externalization process”.¹¹ Someone who cannot yet say what they know can say what it is like, and the metaphor is then worked into an analogy and, later, a model.¹² The book’s narrative thread belongs here. A story is the first explicit form a tacit understanding takes, which is why accounts travel between people as stories before they travel as reports, as chapters 4 and 8 said.

Combination creates “explicit knowledge from explicit knowledge” by “sorting, adding, recategorizing, and recontextualizing”.¹³ It is what a report does with data, what a dashboard does with reports, what a strategy paper does with dashboards. It is also what a language model does, and chapter 13 returns to that.

Internalization moves explicit knowledge back into tacit form. It “bears some similarity to the traditional notion of ‘learning’”, and its engine is action: people internalize what they have done, through “‘experimentation’”, “a process of ‘learning by doing’”.^{14,15} A procedure that has been read is explicit knowledge. A procedure that has been followed a hundred times under varying conditions has become something else, and the person who holds it can now recognize the case the procedure did not cover.

The four together, in Nonaka’s picture, form a spiral. Knowledge “is created by individuals”, and the organization “‘organizationally’ amplifies” it; creation runs as “an

¹¹Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

¹²Nonaka (1994), pp. 20–21 (the sequence metaphor, analogy, model).

¹³Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

¹⁴Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

¹⁵Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

upward spiral process, starting at the individual level moving up to the collective (group) level, and then to the organizational level, sometimes reaching out to the interorganizational level”.^{16,17} The 1986 article had already named the multi-level form. It called it “multilearning”: learning “across multiple levels (individual, group, and corporate) and across multiple functions”.¹⁸ This is the one explicit theory of learning across levels that the book uses, though it is not the only one: Crossan, Lane and White’s 4I framework is the other standard multi-level account, and the book names it without working from it.¹⁹ The theory maps onto the three levels chapter 1 kept in view: the person, the group, the organization and its units.

What the four modes are for

Two cautions before the four modes are put to work.

The first concerns the word *conversion*. Nonaka writes as if tacit knowledge were converted into explicit knowledge and back, and the metaphor of conversion has been contested since. Haridimos Tsoukas argued that tacit knowledge is displayed — manifested — in what we do, and is re-punctuated, that is articulated, through dialogue, but is never captured, translated or converted into explicit form. Stephen Gourlay went further and questioned the conceptual and empirical basis of the theory as a whole.

¹⁶Nonaka (1994), p. 17 (individuals create knowledge; the organization amplifies it; intention, autonomy and fluctuation; Weick 1979 on enactment), p. 18 (breakdowns after Winograd and Flores). The gloss of “breakdown” as the moment a routine stops working and its assumptions become visible is the author’s; Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

¹⁷Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

¹⁸Takeuchi & Nonaka (1986), reprint p. 7 (“multilearning”).

¹⁹Crossan, Lane and White (1999), the 4I framework of organizational learning. Recorded from the review of the process literature, Pettit, Crossan & Vera (2017), p. 482, which gives it with Nonaka and Takeuchi (1995) as the two multi-level frameworks.

Nonaka and Georg von Krogh replied in 2009.^{20,21,22} All three positions have since been read in the original; the pages are in the footnotes. The book does not need to settle it, because it does not use the four modes as a mechanism. It uses them as a diagnostic vocabulary: four kinds of movement that an organization either provides for or does not. Whether what moves is converted, or shown and re-grown in another person, is a question the diagnostic can leave open.

The second caution concerns the sample. The 1986 findings were “drawn, for the most part, from observing how the development process was managed in Japan”, and the authors say so themselves: “General conclusions ... must be made with some caution.”²³ The 1994 theory generalizes them. A reader who has watched the theory travel through knowledge-management programs in organizations that resemble a Japanese manufacturer in nothing will not need the warning, but it belongs in the footnote all the same.

Used as a diagnostic, the four modes give a question that a design can answer: which of the four movements does this organization provide for, and which does it leave to chance? My answer for most organizations I have seen is the same, and it is the chapter’s second claim: strong in combination, weak in the other three. Combination is what an organization’s formal systems do all day. Reports

²⁰Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed., ch. 21). Wiley. (chapter pp. 17–18 of 21): “Tacit knowledge cannot be ‘captured,’ ‘translated,’ or ‘converted’ but displayed—manifested—in what we do.”

²¹Gourlay, S. (2006). Conceptualizing knowledge creation: A critique of Nonaka’s theory. *Journal of Management Studies*, 43(7), 1415–1436. Abstract, p. 1415: “Nonaka’s proposition ... is flawed. Three of the modes appear plausible but none are supported by evidence that cannot be explained more simply”.

²²Nonaka, I., & von Krogh, G. (2009). Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization Science*, 20(3), 635–652. p. 638 responds to Gourlay (2006) by name; p. 637: “Tsoukas (1996, 2003) ... noted that ...”.

²³Takeuchi & Nonaka (1986), reprint p. 9 (institutionalization and unlearning; the limits of the sample).

are written from data; dashboards are built from reports; knowledge bases collect the articles; the strategy paper recategorizes it all. An organization with a good intranet and a documentation standard can believe that it manages knowledge, because it manages the kind of knowledge that is already explicit, and it manages that well. What it loses is the rest: what its people know from experience and have never said, because the pairing and the handover were cut as inefficiencies; what could have been said and never was, because nobody was asked what the last project was like, only what it delivered; and what was written and never practiced, because the procedure was read at onboarding and not since.

The pattern has a mechanism, and it is the one chapter 2 gave for attention. Combination is visible and countable; the other three are not. A knowledge base has an article count. Socialization has no count, and an organization that manages by what it can count will fund the first and starve the other three without ever deciding to. The diagnostic question therefore asks where, in this organization, tacit knowledge becomes shared, where it becomes words, and where words become practice again; whether the organization “manages knowledge” tells one nothing. The answer is usually a person, a place or a habit that no organization chart names and that the next reorganization will remove.

Episode — 4,200 articles and three weeks of leave

In the annual report of the After-Sales Service unit at Halvern, a manufacturer of heat-pump systems for commercial buildings, the knowledge base had a page of its own: about 4,200 articles, each with a named owner and a review date, written to a documentation standard and linked to the service-bulletin system. A new bulletin reached all 180 field technicians on the day it was approved, and the system recorded who had opened it. By every count the unit could make, it managed its knowledge well.

What it had stopped providing for had no page. Until the efficiency program of the year before, the technicians of a region had met at the depot in the morning to collect parts, and for some weeks each year a junior had ridden along with a senior. The program ended both, and technicians now started alone from home. Nobody had decided that they should stop telling each other about last week's jobs; the depot mornings had been booked as travel time. The nearest replacement was a messenger group that technicians in region South had set up themselves.

The loss showed in March of the first winter after the launch of the C4 controller. For some weeks technicians had been applying two settings that Milan, a senior second-line specialist, had worked out for units that started and stopped their compressors far too often. The settings traveled through the messenger group in one message. What did not travel was what Milan did when a technician phoned him: he asked what the building-management system did at night, how many starts the unit made in an hour, and what had been connected when. Then Milan went on leave for three weeks, and the cases in which the two settings did not hold stalled. His colleagues on the second line had the settings. They did not have the cases behind them, and nothing in the 4,200 articles said how he told one unit from another. Nobody had asked him.

In May the official service bulletin was approved, and the knowledge-base article followed the standard: title, affected product, numbered steps. It gave the steps and left out the conditions. That the unit had to sit behind a building-management system overriding its setpoints was something Milan and Tobias, a field technician in the South, knew from the units they had stood in front of, and the template had no field for it. Technicians opened the article, and the system showed it as read. Some applied it to units with no

building-management system at all, which then ran with a longer off-time than they needed, and building operators complained about efficiency. None of those technicians had applied the steps once beside someone who knew when they applied.

The review of the season worked from the ticket counts and the 96 replaced compressors. No technician was asked what a short-cycling unit had been like on site. By the summer the knowledge base held the steps, correctly formatted and owned. How to tell which unit needed them was still held by Milan, by Tobias and by the technicians who had phoned them.

The Palau: a solution found, and not saved

Chapter 1 told the *Palau* story in full, and the second half of it belongs here. When the gyrocompass failed, the navigation team took over a computation the instrument normally performed, and in the twenty-five minutes before the ship was anchored it found a new division of labor under pressure. Hutchins' analysis of that reorganization turned up two findings that this chapter has to hold together with Nonaka.

The first runs against Nonaka's starting point. The new way of dividing the work "was clearly discovered by the organization before it was discovered by any of the participants".²⁴ Nobody designed it; nobody could have said what it was while it was happening; it was "evolutionary in the sense that the system-level change that resulted was never represented".²⁵ In Nonaka and Takeuchi's picture, new knowledge starts with individuals and the organization amplifies it into something shared. On the *Palau* the organization created a solution and no individual held it. The two theories weigh the individual and the group differently, and the book has named the difference since chapter

²⁴Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

²⁵Hutchins (1995), ch. 8, p. 349.

1 without settling it. The most useful reading I can offer is that they describe two different things. Nonaka describes how what an individual knows becomes shared; Hutchins describes where the shared result lives and how it works; and the *Palau* shows that the second can occur without the first.²⁶ That is my reading and neither author's.

The second finding is the warning. "The solution to the problem of organizing work that was discovered by the navigation team was not saved in the system."²⁷ The team learned something under pressure and had no way of keeping it. In the four-mode vocabulary, the *Palau* had a burst of socialization, intense and wordless, under fluctuation of exactly the kind Nonaka names; no externalization, because nobody could say what had happened; no combination, because there was nothing explicit to combine; and no internalization beyond what each person retained of his own part. The spiral did not turn. The organization learned, in Hutchins' strong sense, and then forgot, because what it had learned existed only in the configuration of a team during twenty-five minutes.

Hutchins draws a conclusion that joins this chapter to the next. "The processes by which work is accomplished, by which people are transformed from novices into experts, and by which work practices evolve are all the same processes."²⁸ Doing the work, learning the work and changing the work are one activity looked at three ways. That is a stronger claim than Nonaka's spiral, and a more uncomfortable one for anyone who runs a training department, because it says that learning cannot be scheduled apart from working and then attached to it. He gives the mechanism elsewhere in the book: "every error-correction event is a learning context", and the context extends to everyone who witnesses the correction, and not

²⁶Author's reading. (author's reading)

²⁷Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

²⁸Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

only to the person who made the error.²⁹ The organization that wants to keep what the *Palau* lost has to catch it while the work is being done, in front of the people who will need it next, which is what socialization and externalization are for.

Where the learning goes

If the solution was not saved in the system, where did it go? The question sounds rhetorical and is not. Organizational memory has been studied, and the standard account, James Walsh and Gerardo Ungson's 1991 article, treats it as a set of retention facilities: the memories of individuals, the culture, the transformations (procedures and routines), the structures (roles), the physical setting, and archives outside the organization.³⁰ I have that list from secondary sources. What matters for the chapter is its shape. Memory is distributed across people, artifacts and structure in the same way that understanding was in chapter 7, and each facility forgets in its own way. Individuals leave. Culture drifts. Procedures are rewritten. Roles are reorganized. Buildings are sold. The transactive-memory idea from chapter 7, knowing who knows what, adds a seventh facility, and it is the one most organizations lose first: when the person who knew whom to ask retires, everything she pointed to is still there and none of it can be found.

Hutchins adds the part that Walsh and Ungson's list can hold but does not stress. Artifacts are not only where explicit knowledge is stored. They are where past solutions are built in, so that the present can run on them without knowing they are solutions. Implementations of tasks are "part of a cultural process that tends to collect representations that permit tasks to be performed by means of simple cognitive processes".³¹ The bearing log, the chart,

²⁹Hutchins (1995), ch. 6, p. 279.

³⁰Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91. p. 62, p. 81 (six retention facilities).

³¹Hutchins (1995), ch. 3, p. 155.

the form: each contains the solved problems of earlier generations, and each lets the current team think less because someone before them thought more. "Humans create their cognitive powers by creating the environments in which they exercise those powers."³² An organization's memory is, to a large degree, its tools, and that is why the *Palau's* solution was lost. It never became a tool. Nobody changed the log, the procedure or the chart to hold it, and the next team to lose a gyrocompass would start from where this one had.

Where the learning goes, then, is into whichever facility caught it, and the facilities do not catch the same things. A person catches the tacit; an artifact catches the explicit; a structure catches the division of labor. An organization that does not know which facility holds what has no way to check, when a person leaves or a system is replaced, what it has just forgotten. That check is a design item, and chapter 14 lists it.

Four meanings of "loop"

The chapter's title uses a word that the management literature uses in at least four senses. The book has to keep them apart, because the four are regularly confused and the confusions have consequences. Much of the vocabulary reached practitioners through Peter Senge's *The Fifth Discipline*, which I do not draw on here beyond noting that its "learning organization" folded several of these senses into one phrase.³³

The first sense is Nonaka's spiral: the movement of knowledge through the four modes and up the levels. It is a loop in the sense that internalized knowledge becomes the tacit base for the next round of socialization. Its time scale is months to years.

The second is Chris Argyris and Donald Schön's. In their 1978 book they separated single-loop learning, in which an error is detected and the action corrected within

³²Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

³³Senge, P. M. (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday. (named for the term only).

the organization's existing governing assumptions, from double-loop learning, in which the assumptions themselves are questioned and changed.³⁴ The distinction is theirs, and the summary above follows their own statement of it. In the vocabulary of chapter 5, single-loop learning revises the action and double-loop learning revises the model. The corporate plan that Beer wanted “continuously aborting”, in chapter 1, aborts by double-loop learning; a plan that is adjusted without its assumptions being touched is single-loop, and it can be adjusted forever. Nonaka's view of the matter is sharper than the usual one. He treats double-loop learning as “not a special, difficult task but a daily activity for the organization”.³⁵ An organization that creates knowledge is revising its model all the time, and the drama usually attached to questioning one's assumptions is a sign that the revision has been postponed too long.

The third sense is Boyd's, and the earlier chapters have already met it. Chapter 1 drew the inner circle of the loop, from cue to account to action and back to new cues, and said it can turn fast. Frans Osinga's reconstruction of Boyd's full loop makes the inner circle itself a learning device: decisions “feed back into the systems as validity checks on the correctness and adequacy of the existing orientation patterns”, and, in Osinga's words, “the entire OODA loop is a double-loop learning process”.³⁶ That claim is the bridge between the inner and the outer circle. Every turn of the inner circle tests the model. The outer

³⁴Argyris, C., & Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley. Single-loop learning is defined in the Introduction, pp. 2–3, as the “error-detection-and-correction process” that lets the organization “carry on its present policies or achieve its present objectives”; double-loop learning is defined there too — “Double-loop learning occurs when error is detected and corrected in ways that involve the modification of an organization's underlying norms, policies, and objectives” — and both are developed in ch. 2, “Single-loop learning” (p. 18) and “Double-loop learning” (pp. 20–22).

³⁵Nonaka (1994), p. 19 (the four modes; metaphor and action; double-loop learning as a daily activity).

³⁶Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge, ch. 6, fig. 6.1 “The real OODA loop” and the surrounding text. (the PDF has no page numbers).

circle, chapter 1's slow one, is where the tests accumulate into changed capability. The operating model the book grew out of couples two such loops, one for strategy and one for leadership, and calls the coupling a double loop as well; the term is the same, the thing is different, and Appendix A keeps them apart.³⁷

The fourth sense is the plainest and the one practitioners mean most often: a review rhythm. A retrospective every two weeks, a strategy review every quarter, a portfolio review every year. This is a loop in the calendar, and it is the temporality thread of the book made into a design variable: how often does each level of the organization stop to ask what its action produced? A rhythm by itself learns nothing. It is where the other three loops are given a chance to turn, and an organization without one leaves them to accident. Chapter 14 treats the rhythm as a grid of its own.

Confusing the four costs something specific. An organization that installs a review rhythm and calls it double-loop learning has scheduled the meeting and not the revision. One that runs Boyd's inner circle fast and calls it learning has tested its model many times and changed it never. One that runs the spiral in a knowledge-management program and calls it double-loop learning has moved knowledge between forms without asking whether the assumptions behind the forms were still true. Keeping the senses apart is what lets the design in chapter 14 ask, for each of them, whether it exists.

Unlearning, and why it is skipped

Takeuchi and Nonaka's warning of 1986 can now be given its mechanism. Institutionalization is the point at which what an organization learned stops being held by people who can revise it and starts being held by procedures and artifacts that cannot. "Passing down words of wisdom from the past or establishing standard practices

³⁷The coupled loops of the operating model the book grew out of, and the several meanings of "double loop", are mapped in Appendix A. Standalone readers need nothing from them here.

based on success stories works well when the external environment is stable.”³⁸ The sentence has an unstated second half, which is the epistemic premise of chapter 1 seen from the other side: where the environment changes faster than it can be known, a standard practice based on a success story is a model of a situation that has passed, executed with confidence.

Unlearning is the deliberate discarding of such a practice, and organizations skip it for a reason that Hutchins makes precise. What an organization learned is built into the environment in which it thinks. The form has fields for the things that mattered last time. The approval chain routes decisions to the people who were right last time. The dashboard shows the axes along which the last change moved. None of these announces itself as a past lesson; each presents itself as how things are done. To unlearn is therefore to rebuild part of the environment that does one’s thinking, while thinking with it, and there is no moment in an operating calendar when that is convenient. Chapter 1’s measurement trap is one instance. The metrics encode an old picture, and the organization that reads them is reading its own past understanding as the present.

Beer gives the positive form of the same idea in a sentence about pupils and teachers. When each builds a model of the other and “the two models of each other converge, learning has occurred”.³⁹ Learning, on that reading, is the convergence of a model with what it models. Unlearning is what has to happen when the thing modeled has moved and the model, now built into procedure, has not. The two are one process with two directions, and an organization that funds the first and has no procedure for the second will, over time, hold a model that was once excellent.

There is a further reason unlearning is skipped, and it belongs to chapter 10. A practice based on a success story has an author, and the author is often still in the building.

³⁸Takeuchi & Nonaka (1986), reprint p. 9 (institutionalization and unlearning; the limits of the sample).

³⁹Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

To unlearn it is to say, in the vocabulary of the last chapter, that an account someone gave and the organization committed to no longer counts. Restricted sensemaking protects institutionalized lessons. Guided sensemaking is the form in which they can be questioned, and the leader's forum is where it happens, which is one more reason the forum has to stay open after the leader has formed a view.

What learning tests against

Double-loop learning revises the model, and a model can only be revised against something. Chapter 6 gave the something a name. A decision, in the book's sense, is an explicit and accountable commitment, and a decision record states what was committed to, by whom, and on what expectation. Weick's reading of Garfinkel's jurors said why the record matters: without one, "the person defines retrospectively the decisions that have been made", and the organization rewrites its choices after the outcome is known.⁴⁰ An organization that does that cannot learn, because the expectation the outcome should have tested has been adjusted to fit the outcome. With a record, a decision is Boyd's hypothesis in written form, and the review rhythm is when the hypothesis is compared with what happened.

That gives the four modes their place in the outer circle. The decision record is combination, explicit from explicit. The comparison with the outcome is where fluctuation enters, because the outcome rarely matches the expectation, and the mismatch is the breakdown that makes assumptions visible. Saying what the mismatch was like, before saying what it was, is externalization, and it starts with a story; a review that begins from the metrics and never asks anyone what happened produces combination and nothing else. What the people who were there take away from the review and carry into the next round is internalization. And the pairing of people across rounds,

⁴⁰Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 10–11, on Garfinkel's jurors.

the osmosis of 1986, is socialization. The loop closes only if all four are provided for, and the record is what makes the first move possible.

What would count against it

The chapter claims that an organization which deliberately provides for all four movements of knowledge retains and spreads what it learns better than one that relies on documentation, and that what it does not provide for is lost in the ways the *Palau* and Walsh and Ungson's facilities describe. Against it would count: organizations that run all four modes deliberately not retaining and spreading new knowledge better than organizations that rely on documentation alone; organizations with documentation alone recovering solutions of the *Palau* kind as often as those with socialization and externalization in place. Where to look: evaluations of knowledge-management programs, which have mostly measured combination; Gourlay's critique of the empirical basis of Nonaka's theory, which has to be answered rather than cited around; and the studies of expertise loss on retirement, where the difference between a documented and a practiced process shows most plainly.

Handover: where learning settles

Return to the 1986 article. The companies it described moved what they had learned from project to project by moving people. That was the mechanism, and it worked because what the people carried was tacit and could not have been moved any other way. The *Palau* is the case in which even that mechanism had nothing to work on: the solution belonged to no person, and so no person could carry it to the next ship. Between the two cases lies the design question of chapter 14, which of the four movements an organization provides for, and the question of chapter 12, which is what happens to learning when it does settle in a person. Knowledge that has been internalized does not stay knowledge. Practiced under feedback, it becomes

skill, and in time something the book has been calling judgment since chapter 1: a rich, largely tacit model, tested in practice, that lets a person notice the next cue and recognize the action. Chapter 12 follows it there.

References

- Argyris, C., & Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Crossan, Lane, & White (1999). The 4I framework of organizational learning.
- Gourlay, S. (2006). Conceptualizing knowledge creation: A critique of Nonaka's theory. *Journal of Management Studies*, 43(7), 1415–1436.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37. <https://doi.org/10.1287/orsc.5.1.14>
- Nonaka, I., & Takeuchi, H. (1995). *The Knowledge-Creating Company*. Oxford University Press.
- Nonaka, I., & von Krogh, G. (2009). Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization Science*, 20(3), 635–652.
- Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. Routledge.
- Pettit, Crossan, & Vera (2017).
- Polanyi, M. (1966). *The Tacit Dimension*.
- Senge, P. M. (1990). *The Fifth Discipline: The Art and Practice of the Learning Organization*. Doubleday.
- Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116.
- Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed., ch. 21). Wiley.

- Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16(1), 57–91.
- Weick, K. E. (1979).
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.
- Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition*. Ablex.

Chapter 12

Chapter 12 — The skill development cycle

The simplest part of the task

“In the world of navigation, as in many other systems, novices begin by doing the simplest parts of a collaborative-work task.”¹ Edwin Hutchins wrote that sentence about the navigation team whose work chapter 1 described and chapter 7 traced step by step. The simplest part is reading a bearing off a scale and saying the digits aloud. A newcomer can do it on the first day. They do not yet know what the digits are for, where they go, or what the plotter will do with them at the chart, and they do not need to know, because “much of the structure required for them to organize their activity is available in the social relations”.² Someone tells them when to read. Someone receives what they say. Someone corrects them when the number is wrong. The newcomer is taking part in a computation that nobody has explained to them and that they could not yet carry out alone.

If they stay, they may one day be the one at the chart. Between the two positions lies the subject of this chapter. Chapter 11 asked how an organization learns from what its action produced and why so much of what it learns is lost. This chapter asks about the part that is kept, the part that settles in people. The loop drawn in chapter 1 has an outer circle, from action to learning to changed capability to changed attention, and it turns over years. This is where it turns.

¹Hutchins (1995), ch. 6, p. 264.

²Hutchins (1995), ch. 4, p. 226.

Five capabilities

Learning feeds capabilities. The word covers several things that a job profile treats as one and that come apart as soon as someone has to grow them in a person. The book works with five: knowledge, skill, proficiency, fluency and judgment.³ The list and its order are settled — knowledge, skill, proficiency, fluency, judgment — and the cycle they form is read as a design rather than a description; what the five refer to, and the cycle's own order and name, can be stated now.

Knowledge is what can be said and taught: that bearings are read from this scale in this way, that a filing is due within thirty days, that customers of a certain kind tend to leave after their second renewal. It can be written into a manual, delivered in a room and checked with a question.

Skill is the ability to do a task to a standard, acquired by doing it. The newcomer who has read bearings for a week has a skill they did not have on the first day, and no amount of reading about bearings would have given it to them. Skill can be shown and assessed.

Proficiency is skill that holds up under the ordinary variation of conditions: bad light, a moving deck, a distracted recorder, a customer who does not follow the script. The proficient person does the task reliably without someone standing behind them.

Fluency is performance that no longer takes the person's attention. They can do the task while attending to something else, and they can vary it without stopping to think. The plotter who talks to the navigator while drawing the line is fluent.

Judgment is knowing what the situation calls for when the procedure does not say, including when the procedure should be set aside and when nothing should be done. It is what the *Palau's* team had to draw on when the instrument failed and the procedure with it, in whatever form the team

³The five capabilities follow the "Learning — SDLC → Capabilities" drawing of the Tautai operating model and its fluency levels (Appendix A). The definitions given here are the author's working definitions.

held it, and it is what the Mann Gulch crew could not build in the minutes they had.

Two things about the list are visible already. The five differ in kind: an organization can be full of knowledge and short of skill, or fluent in a task that no longer needs doing, or proficient across the board and without anyone who knows when to stop. And each of the five is present on the navigation team at once, distributed over its roles, from the newcomer at the scale to the navigator who decides what the fix means. Organizations rarely keep the five apart. A job profile asks for “experience”, counted in years, and a career ladder marks the years with titles; both assume that time in a role produces the capabilities the role needs, and neither says which of the five it expects. Time in a role produces some of them under some conditions, and the rest of this chapter is about the conditions. Chapter 11 read learning through Nonaka’s theory of knowledge creation. The same theory sorts the five.

The five read with SECI

Ikujiro Nonaka, following Michael Polanyi’s “We can know more than we can tell”, separated explicit knowledge, “transmittable in formal, systematic language”, from tacit knowledge, which “has a personal quality, which makes it hard to formalize and communicate” and is “deeply rooted in action, commitment, and involvement in a specific context”.⁴ He gave tacit knowledge two components: a cognitive one, the mental models a person carries, and a technical one, “concrete know-how, crafts, and skills”.⁵

Put the five capabilities beside that distinction and they sort themselves by how much of each is tacit. Knowledge sits at the explicit end; it is what can be told. From skill

⁴Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4 (public PDF, pp. 14–21).

⁵Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 16, citing Polanyi, M. (1966), *The Tacit Dimension*, p. 4 (public PDF, pp. 14–21).

onward the tacit share grows. Skill, proficiency and fluency are Nonaka's technical component: know-how that shows in performance and resists being said. Judgment is the cognitive component: the model of the situation, of what matters in it and of what it will do next, which the person cannot state in full and which they act on anyway. This reading is mine, and it is a reading of the list, not a claim about Nonaka; what it gives the chapter is a single axis along which the five differ, and a prediction about how each of them can be passed on.

Nonaka's four modes of conversion supply the prediction.⁶ *Socialization* creates "tacit knowledge through shared experience"; his example is the apprentice, who learns "by observation, imitation, and practice". *Internalization* turns explicit knowledge into tacit and "bears some similarity to the traditional notion of 'learning'"; its mechanism is "'experimentation' ... a process of 'learning by doing'".⁷ *Combination* makes "explicit knowledge from explicit knowledge" by "sorting, adding, recategorizing, and recontextualizing": the course, the manual, the knowledge base, the training video. *Externalization* is the attempt to say what one knows tacitly, and Nonaka puts metaphor at its start.

Read along the axis, the modes divide. Knowledge travels by combination; that is what combination is for. Skill, proficiency and fluency are grown by socialization and internalization, by working beside someone and by doing the task under conditions that answer back. Judgment is grown the same way and takes the longest, because the model it consists of has to be built from many situations, and situations arrive at their own pace. Combination does not produce it. A course can give a person the words for what an expert does; it cannot give them the expert's recognition of the situation in which those words apply.

Externalization sits across the axis rather than along it. When an experienced person tries to say what they do,

⁶Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

⁷Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

Nonaka observes, the first step is usually a metaphor, and the sequence runs from metaphor through analogy to a model that others can use.⁸ What comes out is knowledge for the next person, at the explicit end of the axis, and it is worth less than what went in: the metaphor carries the shape of the expert's model, and the recognition that would apply it stays behind. This is not a fault in the expert's articulation. It is what the tacit dimension means. The organization gains a teachable account, and the account is a starting point for someone else's trajectory, not a substitute for it.

Chapter 11 attached a caution to the four modes, and it applies here with more force. It is contested whether tacit knowledge is "converted" into anything. Tsoukas argues that tacit knowledge is articulated through dialogue but is never converted or translated into explicit form.⁹ The book uses the modes as a vocabulary for asking which occasions for learning an organization provides, not as a claim about what happens inside a person. Asked that way, the diagnostic is short. An organization with a training catalog, a documentation standard and a knowledge base, and no arrangement by which a newcomer works beside someone who can already do the job, runs one mode and expects it to deliver all five capabilities. It delivers the first.

Episode — The Second-Line Academy

"Outside tree. Forwarded to senior." In the winter after the first cohort was certified, that note closed a good share of the escalations handled by the newest technical specialists in the after-sales service of Halvern, a manufacturer of heat-pump systems for commercial

⁸Nonaka (1994), pp. 20–21 (metaphor as the first step of externalization). *The Tautai Principle*, ch. 12 "Steering with Stories" as listed in the abstract.

⁹Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed., ch. 21). Wiley. (chapter pp. 17–18 of 21): "Tacit knowledge cannot be 'captured,' 'translated,' or 'converted' but displayed—manifested—in what we do."

buildings. The seniors it referred to were three people. Milan, nineteen years in the company, was one of them.

The autumn before, the second line had four openings and no internal candidate. Specialists had always come from the field after five to eight years as technicians, and Katrin, the Head of Service, did not have five years. Service management built the Second-Line Academy: nine e-learning modules, three classroom days, a set of decision trees for the common fault families, and a certificate. The material was drawn from the knowledge base and the service bulletins, and it was sound. Milan reviewed the decision trees himself and corrected two. Eight people were certified in the first cohort and began taking escalations.

They knew the vocabulary. Asked what a parameter did, they could say. Given a case the trees covered, they worked it branch by branch to the right leaf, somewhat slower than the seniors and with the same result. Then an operator reported a unit that tripped on high pressure on mild afternoons and ran without fault in the cold. One of the graduates checked refrigerant charge, fan, sensor values and firmware version. Each came back within limits, the tree ended, and she wrote the note. Milan read the ticket and asked the operator a question that stood on no tree: what else had changed in the building that month. A new ventilation outlet had been fitted two meters from the outdoor unit. The ticket returned to the graduate closed, with cause and remedy in two lines. Why Milan had asked about the building was not in it.

Katrin reviewed the Academy after the heating season. On cases inside the trees the graduates' work held up. The share of cases they forwarded had not fallen between November and February, and the three seniors' queues were longer than they had been before the openings were filled. Asked in the review what the graduates lacked, Milan said they knew what every

parameter meant and had not yet seen enough units to know which one to look at.

The review produced a tenth module, on atypical cases. It also recorded how the graduates had worked through the winter: each at a separate desk or at home, taking cases from the queue and handing them on through the ticket system. In four months none of them had sat beside a senior during a diagnosis, and the ticket system showed them the result of a forwarded case and nothing of how it had been reached.

Trajectories within a horizon of observation

Hutchins found the same thing from the other side, watching how the navigation team turned newcomers into navigators. The mechanism has three parts in his account, and none of them is instruction.

The first is the trajectory of roles. Newcomers begin with the simplest parts of the shared task and move, over time, through positions that carry more of it. The sequence is set by the work, not by a curriculum: the roles exist because the task needs them, and the fact that they can be occupied in order of difficulty is a property of the system that the learner benefits from. Nonaka's apprentice and Hutchins' novice are the same figure seen from two sides. Nonaka describes what the individual acquires; Hutchins describes what the system provides for them to acquire it in.

The second is the horizon of observation: what a person in one position can see and hear of the work done in the others. The newcomer at the scale learns more if they can see where their digits go than if they cannot. And that horizon is set by the arrangement of the work rather than by the person's curiosity. Hutchins observed that "the design of a tool may change the horizon of observation of those in the vicinity of the tool".¹⁰ A chart laid out where the team can see it opens the horizon; an instrument that

¹⁰Hutchins (1995), ch. 6, p. 270.

computes the same result inside a housing closes it. Every artifact in the system therefore decides, as a side effect of its design, who can learn what by watching.

The third is error. In Hutchins' words, "every error-correction event is a learning context not just for the person who commits the error but for all who witness it".¹¹ On the navigation team, errors are caught by the next person in the chain, aloud, and the correction is heard by everyone within the horizon. That is how the newcomer learns what a wrong bearing sounds like before they are the one who has to catch it. It is also why the team's overlapping knowledge, which chapter 8 described as "an important source of the robustness of such systems in the face of error and interruption",¹² and its capacity to grow people are one arrangement rather than two. The overlap exists because people have passed through each other's roles; the robustness exists because of the overlap; and the passing through is the trajectory. Hutchins drew the conclusion at the end of his analysis of the *Palau's* emergency: "The processes by which work is accomplished, by which people are transformed from novices into experts, and by which work practices evolve are all the same processes."¹³

Hirotaka Takeuchi and Nonaka had seen a version of this in product development a decade earlier: the transfer of learning from one project to the next happened "through 'osmosis'", by assigning key people to the next project, rather than through documents.¹⁴ The trajectory there runs across projects instead of across roles, and the mechanism is the same: people carry what they know into a new shared task, and others acquire it by working beside them.

The expertise literature adds a detail that will matter for the machine question. Experts, Reep's review finds,

¹¹Hutchins (1995), ch. 6, p. 279.

¹²Hutchins (1995), ch. 4, pp. 221 and 227.

¹³Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

¹⁴Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116, p. 9 (institutionalization and unlearning).

distribute their cognition into their context and their artifacts; novices cannot yet do this.¹⁵ The expert's judgment is more than a richer model in the head: it includes a practiced way of using the environment, of knowing which instrument to trust, which record to check and which colleague to ask. Hutchins put the general form of this as "Humans create their cognitive powers by creating the environments in which they exercise those powers."¹⁶ The newcomer has to learn the environment as well as the task, and can only do so from inside it.

For an organization the three parts translate without much loss. Are there entry roles from which the shared task can be seen? Do people move through roles in an order that carries more of the task each time, or do they arrive in a role and stay? Can a junior person see the work of the senior one, or do the tools, the meeting formats and the office layout hide it? Are errors corrected where others can witness the correction, or in private, or not at all? Every one of these is a design decision, and most organizations make it by accident.

Take a support function as an illustration. Where first-line agents handle the simple tickets and can see how the second line resolves the ones they escalate, the queue is a trajectory with an open horizon: the agent learns what a hard case looks like from the cases they could not solve. Where the escalation disappears into another team's system and comes back as a closed ticket, the same agent does the same simple work for years and learns nothing from the hard cases that passed through their hands. The two arrangements cost about the same. They grow different people.

¹⁵Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 8, pp. 187–191). CRC Press.

¹⁶Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

A cycle from what can be taught to what can only be grown

Several traditions have described the path from newcomer to expert, and the book relates its five capabilities to two of them.

Hubert and Stuart Dreyfus set out five stages of skill acquisition, from novice to expert.¹⁷ In their account the novice applies rules to features that have been defined for them without regard to context; with experience the person begins to see situations rather than features, and the expert responds to the situation as a whole without deliberating and often without being able to say what rule they followed. The Japanese teaching tradition compresses a similar path into three words, *shu*, *ha*, *ri*: keep the form, break the form, leave the form.¹⁸ Both describe a movement from what can be told to what can only be done, and both end where rules run out and judgment begins. Anders Ericsson's work on expertise supplies the condition the movement depends on: expert performance comes from practice arranged to improve it, with feedback, sustained over years; experience as such, the same task repeated without feedback, does not produce it.¹⁹ Gary Klein's finding from chapter 6 fits the same picture from the decision side: comparing options was what novices did; the experienced commanders recognized the situation and the action came with it.²⁰

The chapter's proposal is to read the five capabilities as a cycle rather than a ladder. Knowledge is where an

¹⁷Dreyfus, S. E., & Dreyfus, H. L. (1980). *A Five-Stage Model of the Mental Activities Involved in Directed Skill Acquisition*. Operations Research Center, University of California, Berkeley; Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over Machine*. Free Press. (the 1980 report's stages are "novice, competence, proficiency, expertise and mastery"; the "novice → expert" labeling belongs to the 1986 book, *Mind over Machine* — cite dreyfus1986 for the stage names).

¹⁸*Shu-ha-ri*, the three-stage account of learning in Japanese martial and craft traditions.

¹⁹Ericsson, K. A., on deliberate practice and expert performance; the work to cite is still to be specified.

²⁰Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press, pp. 12, 17 and 24 (printed page = PDF page – 24).

organization can start with a person: it can be taught, and combination delivers it. Skill, proficiency and fluency are grown, through socialization and internalization, in a trajectory of roles. Judgment is grown last and slowest. And then the movement turns back. Judgment is what lets a person say some of what they know, which is externalization, and what lets them decide what the next newcomer should be taught first, which is knowledge again. An organization's stock of teachable knowledge is, in this reading, a by-product of the judgment of the people who could put it into words. The cycle is closed at the top, and it runs through people at different points on it at the same time.

The order is settled for the five as a ladder: proficiency precedes fluency, and judgment comes after fluency, grown last and slowest. The order and the name of the cycle are settled as well. On the order, the cycle's order is the five-step ladder; how it is entered and drawn is a matter of emphasis rather than a decision — it can be drawn as a ladder that turns back only at judgment, or read from the first day as a loop the newcomer is already inside, as Hutchins' newcomer, who witnesses error corrections before they can make the error, suggests. On the name, the cycle is read as a design — what an organization should arrange — rather than a description of what happens to a person, the design reading this chapter and chapter 14 use. The Dreyfus stages describe the individual and say nothing about the system; Hutchins describes the system and treats the individual's growth as its output. The chapter needs both, and it keeps the disagreement between Nonaka and Hutchins about who the knower is where chapter 11 left it, open. The Dreyfus stages describe the individual and say nothing about the system; Hutchins describes the system and treats the individual's growth as its output. The chapter needs both. The individual grows the judgment. The system provides the trajectory along which it can be grown.

Closing the loop: judgment as a practice-tested internal model

Chapter 1 named the internal model as the account an organization acts on, held in heads, artifacts and structure. Chapter 5 developed it as what orientation keeps alive, and chapter 6 showed it at work in Klein's commanders, for whom recognizing the situation as a type delivered the action with it. Judgment, in the terms the book has been using, is that model in its individual form: a rich, practice-tested, largely tacit internal model.

Each of the three adjectives does work. *Rich* means it makes many distinctions. Ross Ashby noted that a set's variety is not a property of the set alone; "the observer and his powers of discrimination may have to be specified".²¹ The expert sees states in a situation that the newcomer cannot see, and chapter 2 argued that this is what a weak signal depends on: a cue that fits no distinction in the model is invisible, whatever its strength, and the distinctions are what judgment consists of. *Practice-tested* means the model was built by acting on it and being answered. Stafford Beer described learning between a pupil and a teacher as each building a model of the other until "the two models of each other converge, learning has occurred".²² A model that was only taught has never been tested; a model grown in a trajectory of roles has been corrected at every error-correction event the person witnessed. *Largely tacit* means it cannot be handed over in a document, which is why it takes years and why it is the capability an organization can least afford to lose.

The model also has to be run, and running it is part of judgment. Klein found that his experienced commanders, where no action came with the recognition, tested a candidate by mental simulation, "the ability to imagine people and objects consciously and to transform those people and

²¹Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

²²Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

objects through several transitions”.²³ Chapter 5 described this as the same model run forward. A newcomer can be given the model’s vocabulary; what they cannot be given is the stock of situations against which the simulation is checked, and that stock is what the years in the trajectory supply.

Judgment has a collective form as well, and the book’s two theories of Part IV pull apart on it. For Nonaka, knowledge “is created by individuals” and the organization amplifies it;²⁴ on that reading judgment lives in people and the organization’s share is the trajectory. Hutchins’ *Palau* points the other way: chapter 11 recalled that the team’s new division of labor was found by the organization before any member understood it. Between the two lies what the navigation team plainly had: overlapping knowledge, so that when a role fails the person next to it can take it over, and a role structure practiced often enough that the team can “reconfigure dynamically” under load.²⁵ Whether one calls that collective judgment or a well-designed system for individual judgment to act in, it is grown the same way, by people passing through each other’s roles over time, and it is lost the same way, when they stop.

This is where the loop closes. The book’s argument in chapter 1 was that understanding comes first, and that the five activities form a loop rather than a pipeline. The outer circle of that loop runs from action to learning to capability and back to attention. Judgment is the capability in question. It is what lets people notice the next cues, because noticing depends on the distinctions the model can make (chapter 2). It is what lets them make sense of those cues when the picture breaks, because Weick’s soldiers with the wrong map were saved by the fact that they “kept noticing cues, and ... kept updating their sense of

²³Klein (1998), ch. 5 “The Power of Mental Simulation” (from p. 49).

²⁴Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 17 (knowledge created by individuals; intention, autonomy, fluctuation), p. 18 (the “breakdowns” passage after Winograd and Flores). (public PDF, pp. 14–21).

²⁵Hutchins (1995), ch. 4, pp. 221 and 227.

where they were”,²⁶ and noticing and updating are things a person with judgment does and a person without it cannot. And it is what lets them recognize the action in the situation without a separate act of choosing (chapter 6). Beer’s sentence that “the regulator is as good as its contained model of the system regulated, and no more”²⁷ was written about control systems, and chapter 1 said why the book does not treat it as proved for organizations. In the loose sense the book does use, it can be read here as a statement about people: an organization’s steering is no better than the judgment of the people who steer, and that judgment is grown, not installed.

The temporality of this matters for design. The inner circle of the loop, from cue to account to action, can turn in minutes, and chapter 6 argued that its speed is a consequence of how good the model is. The outer circle turns over years, because that is how long a rich model takes to grow in a person. Chapter 1 noted that organizations often invest in the inner circle and neglect the outer one, and get quicker at acting on a picture they no longer improve. The skill development cycle is the outer circle. Protecting it is a design task, and chapter 14 takes it up under the heading of protected trajectories.

AI removes the entry roles

The chapter has so far treated the trajectory of roles as something an organization has, by design or by accident. The machine thread of the book changes that. When a model does the simplest parts of the task, the newcomer has nowhere to begin.

The Tautai Principle names this the experience-pipeline paradox: the tasks that automation removes first are the entry-level ones, and those are the tasks in which the next

²⁶Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, pp. 54–55.

²⁷Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

generation of experts would have started.²⁸ Hutchins' account turns the paradox into a mechanism. The entry roles are the first steps of the trajectory along which judgment is grown; remove them and the trajectory has no start. In Nonaka's terms, what is removed is the occasion for socialization, the shared experience in which apprentices learn "by observation, imitation, and practice", and with it much of the occasion for internalization, the learning by doing. What remains is combination, which is the mode the machine performs anyway.

The horizon of observation narrows at the same time. A tool that returns a result hides the work that produced it, and it hides it from the people who would have learned by watching. The newcomer who reads bearings sees the chart. The newcomer who receives a finished analysis from a system sees the analysis. Error-correction events go the same way: when the model's output is wrong, whoever is left to catch it does so alone, at a screen, and nobody witnesses the correction. The three parts of Hutchins' mechanism are all affected, and in the same direction.

None of this follows from the machine being bad at the task. It follows from the machine being good at it. The position where the newcomer read digits is now where the model sits, and the model does not need a trajectory. The people do. An organization that lets this happen without noticing becomes, in the phrase chapter 13 will argue for, better informed and less capable.

Episode — The commissioning app and the four openings

In the autumn after the expensive winter, Halvern's After-Sales Service posted four openings for second-line specialists. The posting carried the requirement it had always carried: five to eight years in the field, commissioning included. Katrin, the Head of Service, received three internal applications. None of the three applicants had ever commissioned a system.

²⁸Mathis, K. (2025). *The Tautai Principle*, pp. 199–200 (the experience pipeline). Paraphrased.

Among technicians of that seniority, nobody had. Seven years earlier Halvern had introduced a guided commissioning app, and installer partners had commissioned new systems themselves ever since. Until then commissioning had been the entry role. A junior technician spent the first two years driving to new installations with a senior, reading pressures and temperatures off the controller and calling them out while the senior set the parameters. The junior saw what the senior did with a value that was off and heard the reason. A wrong setting entered by the junior was caught at start-up, on site, with the installer standing by. Milan had begun that way nineteen years before. The app walked the installer through the same steps on a tablet, checked each value against a permitted range and showed a green mark. It did this well, and complaints about commissioning did not rise. The junior commissioning roles were cut and the saving was reported in technician hours. The proposal did not list those roles as the first step toward the second line, because no document had ever described them as that; the posting's two words, "commissioning included", were the only place where the connection was written down. Technicians hired afterwards started with repairs. Their first contact with a heat-pump system was a fault in it, and they did not see one being set up to run. From the year before the C4 launch they also started from home and worked alone, so that what a new technician could observe of a senior's work was the closed ticket.

Katrin went through the list of 180 technicians with the six regional leads. Those who had commissioned beside a senior now had twelve or more years in the field. The ones among them who had wanted the second line were already on it, and the others did not want to move to headquarters. Below them were seven hiring years of technicians who repaired reliably and had not once been present when a senior decided what a

system needed before it ran. The second line counted 22 specialists, seven of them due to retire within five years, and the four openings were the first of those places to come up.

The posting stayed open for two months and was then withdrawn. The Second-Line Academy described earlier in this chapter, with its nine modules, three classroom days and decision trees, was what Halvern put in the place of the two years beside a senior.

What the organization can do about it belongs to the next chapter, which gathers the machine thread into one argument, and to chapter 14, which turns it into design. The question this chapter hands on is narrow: which of the four modes a machine can take over, and which must be kept, by design, as occasions for people.

What would count against it

The chapter claims that judgment is grown along a trajectory of roles inside a working system, through socialization and internalization under feedback, and cannot be produced by instruction and combination alone. It would count against this if judgment were reliably produced by courses, documentation and knowledge bases without practice under feedback; or if novices reached expert performance without a trajectory of roles and a horizon of observation. It would also count against the chapter's "both" if the individual accounts of expertise, Dreyfus's and Ericsson's, turned out to explain expert performance fully without any reference to the system the expert grew in, which would make Hutchins' mechanism redundant. Where to look: the expertise literature, including Ericsson's studies and Reep's review; Klein's naturalistic decision research on what experienced people do that novices do not; Hutchins' sixth chapter on how the navigation team grows its members.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over Machine*. Free Press.
- Dreyfus, S. E., & Dreyfus, H. L. (1980). *A Five-Stage Model of the Mental Activities Involved in Directed Skill Acquisition*. Operations Research Center, University of California, Berkeley.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Mathis, K. (2025). *The Tautai Principle*.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Polanyi, M. (1966). *The Tacit Dimension*.
- Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.
- Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*.
- Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed.). Wiley.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.

Chapter 13

Chapter 13 — Making sense with machines

An amplifier that is not one

“It has now become commonplace to speak of technology ... as an amplifier of cognitive abilities.”¹ Edwin Hutchins wrote that in 1995, about charts, instruments and the other tools of a navigation team, and he wrote it in order to deny it. Following Cole and Griffin, he argued that the appearance of amplification comes from looking at the product of the work instead of the process: “None of the component cognitive abilities has been amplified by the use of any of the tools. Rather, each tool presents the task to the user as a different sort of cognitive problem requiring a different set of cognitive abilities”.² A person with a tool forms a local functional system whose cognitive properties are, in his words, “radically different from the cognitive properties of the person alone”.³ The tool has changed what the person has to be able to do.

Thirty years later the word is back, attached to a new kind of tool, and the chapter takes Hutchins’ sentence as its test. A language model in an organization’s loop is described as amplifying the people who use it. What does it change their task into, and for whom?

Chapter 1 gave the book a first answer without naming it. The *Palau*’s gyrocompass performed a computation

¹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

²Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, pp. 153–154; Cole, M., & Griffin, P. (1980), cited there.

³Hutchins (1995), introduction (chapter summaries; the local functional system).

that nobody on the bridge had to do by hand, and when it stopped the navigation team had to perform that computation by hand, under time pressure, with a division of labor nobody had planned. The instrument had not made anyone better at bearings. It had changed which computation the people had to be able to do, and the change became visible only when the instrument failed.

The thread, gathered

Machines have appeared in four earlier chapters, each time as a participant in one activity of the loop, and the four appearances have a common form.

In chapter 2 they appeared on the input side. A dashboard or a forecasting model encodes the picture that was current when it was built, and it reports movement along the axes that picture expected; pattern detection at scale changes the task of noticing from finding the pattern to deciding which of many found patterns matters, and it does not remove the task. In chapter 6 they appeared on the output side, as selectors: a model that proposes or ranks actions changes the selection task from generating a candidate to choosing among candidates it did not generate. In chapter 7 they appeared as a new medium in the chain of representations, and the question Hutchins' method poses for any medium was posed for them: where does the machine sit in the chain, what does it receive, what does it hand on, and can the next person in the chain inspect and challenge what it hands on? In chapter 12 they appeared at the start of the trajectory of roles, doing the simplest parts of the shared task and thereby removing the positions in which newcomers used to begin.

The common form is this. In each case the machine takes over a part of the work that people did, and in each case what it leaves for the people is a different kind of work, and the same kind: evaluation. Which of the found patterns matters. Which of the ranked candidates fits. Whether the representation handed on is right. Whether

the finished analysis that arrived in place of the newcomer's first draft can be trusted. The machine does not reduce the organization's need for understanding. It moves that need downstream of its own output, and it concentrates it on the people who receive that output.

Abundant content, scarce validated judgment

Herbert Simon's sentence from chapter 1, that "a wealth of information creates a poverty of attention",⁴ was written in 1971 about a world in which producing information still cost something. A model that produces text, code, summaries, analyses and lists of options at a marginal cost close to zero has removed most of that cost. Content is now abundant in the strict sense that its supply no longer limits anything. Simon also gave the design rule that sorts the tools which truly conserve attention from those that do not: a new information-processing subsystem reduces the net demand on the rest of the organization's attention "only if it absorbs more information previously received by others than it produces—that is, if it listens and thinks more than it speaks".⁵

What it does not supply is the thing that makes content usable. Nonaka defined knowledge as "justified true belief", with the weight on "a dynamic human process of justifying personal beliefs".⁶ A model's output is, at best, a candidate belief. The justification — the judgment, in Nonaka and Takeuchi's sense, of whether the candidate belief

⁴Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press. The poverty-of-attention sentence and the design rule are in the OCR'd scan (printed pp. ~40–42); see chapter 1.

⁵Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press. The poverty-of-attention sentence and the design rule are in the OCR'd scan (printed pp. ~40–42); see chapter 1.

⁶Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 15 (public PDF, pp. 14–21).

is warranted against the organization's situation and purpose, not by the situation alone — has to come from somewhere, and it comes from the capability chapter 12 called judgment: a rich, practice-tested, largely tacit model of the situation, grown in people over years. That capability was scarce before the machines arrived. Their arrival makes it scarcer in relation to what it has to be applied to, because the volume of candidates has gone up and the number of people who can tell a good one from a plausible one has not.

Plausibility deserves a word here, because it is what these systems are built to produce. Chapter 4 followed Weick in treating sensemaking as “driven by plausibility rather than accuracy”,⁷ and chapter 1 told of the soldiers with the wrong map who found their way because they kept moving and kept noticing cues. Plausibility was enough to start with because the correction came from the world, through people who noticed. A model's output arrives already plausible; that is its native condition. The noticing that corrects it has to be supplied by the people who read it, and those people need a model of their own against which the output can fail. Where they have none, plausibility is the end of the process rather than its beginning.

The result is a scarcity of a particular kind. Judgment in general was always scarce. What the machines make scarce is validated judgment: judgment that has been tested against outcomes often enough that the organization can rely on it to check what the machine returns. An organization can find out how much of that it has by asking a simple question of any role that receives model output: could the person in this role tell if the output were wrong, and how would they know?

Which conversion modes machines can take over

Nonaka's four modes give the sharpest way to ask what a machine can do for an organization's learning, because

⁷Weick, K. E. (1995). *Sensemaking in Organizations*. Sage; plausibility rather than accuracy as a property of sensemaking.

they were defined by the kind of knowledge that goes in and the kind that comes out.⁸ The book uses them, as chapters 11 and 12 did, as a vocabulary for occasions rather than as a mechanism inside anyone's head. Asked mode by mode, the answer divides cleanly.

Combination creates "explicit knowledge from explicit knowledge" by "sorting, adding, recategorizing, and recontextualizing".⁹ That description, written in 1994, before any of the current systems existed, is a fair description of what a language model does: it takes explicit material and returns explicit material, rearranged, summarized, extended and placed in a new context. This is the mode machines take over well, and it is the mode chapter 11 found most organizations were already strongest in. A machine given to such an organization strengthens the mode it least needed strengthened.

Externalization is the attempt to put what one knows tacitly into a form others can use, and Nonaka placed metaphor at its start.¹⁰ Here a machine can help: it can propose a metaphor, draft the words, offer the structure that an expert then corrects. The risk is in the correction. If the expert reads the draft and finds it plausible, the articulation that would have made the knowledge theirs, and shareable, has not happened. Tsoukas holds that tacit knowledge is articulated through dialogue but is never converted or translated into explicit form;¹¹ whatever one makes of the stronger claim, the weaker one is enough here. Externalization works through the person's own effort to say what they know, and a system that supplies the

⁸Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

⁹Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

¹⁰Nonaka (1994), pp. 20–21 (metaphor as the first step of externalization). *The Tautai Principle*, ch. 12 "Steering with Stories" as listed in the abstract.

¹¹Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed., ch. 21). Wiley. (chapter pp. 17–18 of 21): "Tacit knowledge cannot be 'captured,' 'translated,' or 'converted' but displayed—manifested—in what we do."

saying removes the effort along with the difficulty. Machines can support this mode, on the condition that the person does the work of articulation and the machine does the work of drafting, and not the other way round.

Socialization creates “tacit knowledge through shared experience”, and Nonaka’s example is the apprentice, who learns “by observation, imitation, and practice”.¹² The shared experience is of the task, between people doing it together, and it is what chapter 12 traced along the trajectory of roles: the newcomer at the scale who hears the correction of a wrong bearing. Whatever a model has, it has not done the organization’s task beside the newcomer, and it cannot be the senior colleague whose work the newcomer watches. A machine cannot perform socialization for people. It can crowd it out, by taking the tasks in which it used to happen.

Internalization is learning by doing, “‘experimentation’ ... a process of ‘learning by doing’”.¹³ A machine that does the task for a person has removed the doing. It may have done the task well; the person has still not learned it. Internalization, like socialization, cannot be performed on someone’s behalf, and it is displaced by exactly the delegation that makes the machine useful.

Used as a diagnostic, the four modes give an organization a way to take stock of its own machine use that does not depend on the technology. For each place where a model has been put to work, ask which mode it is doing. Where the answer is combination, the question is only whether the output is checked. Where the answer is externalization, the question is who did the articulating. Where the answer is that the model has taken over a task in which people used to learn by watching or by doing, no mode has been taken over at all; an occasion has been removed, and the inventory should record it as a loss on the learning side, whatever it shows as a gain on the output side.

¹²Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

¹³Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

So: combination, yes; externalization, with care and with the effort left where it belongs; socialization and internalization, no, and at risk. The division is not a verdict on the machines. It follows from what the two tacit modes are: forms of learning that consist in a person's own participation in the work, which is the one thing a machine's participation replaces.

Better informed, less capable

Put the division beside chapter 12's trajectory and the argument of the chapter follows. An organization that lets a machine take over combination gains something real: faster reports, better-organized documentation, summaries of what it already knew. An organization that also lets the machine take the entry-level tasks, the simplest parts of the shared work, removes the start of the trajectory along which its judgment was grown, and with it the occasions for socialization and internalization that the trajectory provided. *The Tautai Principle* named this the experience-pipeline paradox: the tasks that automation removes first are the tasks in which the next generation of experts would have begun.¹⁴ If the organization does nothing to protect those occasions, its stock of explicit knowledge rises and its stock of validated judgment stops being replenished. It is better informed. It is less capable. And the two movements are caused by the same decision, so that the second is easy to miss behind the first.

Three earlier findings say why the loss is hard to see and slow to show.

The first is Hutchins on artifacts. Tools, he found, are "part of a cultural process that tends to collect representations that permit tasks to be performed by means of simple cognitive processes".¹⁵ A model is such an artifact at a new scale: a collection of the explicit knowledge of the past, arranged so that a task which once needed an expert can be performed by someone who asks. Chapter 11 noted, with

¹⁴Mathis, K. (2025). *The Tautai Principle*, pp. 199–200 (the experience pipeline). Paraphrased.

¹⁵Hutchins (1995), ch. 3, p. 155.

Takeuchi and Nonaka, that institutionalized knowledge of this kind “works well when the external environment is stable”¹⁶ and becomes a liability when it changes. A model trained on what was known is institutionalization without the meeting. It carries the same inertia, and chapter 2’s warning about dashboards applies to it: it encodes a picture that was current when it was built.

The second is Reep on experts. Experts distribute their cognition into context and artifacts; novices cannot yet.¹⁷ A model in the hands of an expert is one more artifact into which a practiced judgment can be distributed: the expert knows what to ask, what to distrust and where the output stops being reliable. The same model in the hands of a novice is an oracle, because the novice has no judgment to distribute and no basis on which to distrust. The tool is the same. The functional system it forms with the person is, as Hutchins said, radically different in the two cases. An organization that measures the tool’s value with its experts and then removes the roles in which experts were made has measured the wrong system.

The third is Klein on comparison. Chapter 6 reported his finding that experienced commanders recognized the situation and the action came with it, while comparing options was what novices did.¹⁸ A model that returns a ranked list of options invites everyone who receives it into comparative evaluation, whatever their experience. This reading is mine, and it cuts two ways. For the novice, comparison among options they did not generate is a poor substitute for the recognition they have not yet grown. For the expert, it is a change of mode: the model that used to deliver the action now has to be used to judge a list, and the expert’s recognition becomes one input among several

¹⁶Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116, p. 9 (institutionalization and unlearning).

¹⁷Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In M. D. McNeese, E. Salas & M. R. Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition* (ch. 8, pp. 187–191). CRC Press.

¹⁸Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press, pp. 12, 17 and 24 (printed page = PDF page – 24).

rather than the thing that selects. The selection question of chapter 6 returns in a new form, and the next section takes it up.

Episode — The assistant's summary

“High-pressure lockout, three occurrences in ten days, reset on site each time. Pattern consistent with restricted water flow. Suggested action: check strainer and circulation pump, see linked article.” That was the summary at the top of an escalation that reached the second line of Halvern’s After-Sales Service late in the third year, about a C4 system in a school. Every sentence in it was accurate. A specialist from the first Academy cohort read it, found that the decision tree for high-pressure lockouts led to the same article, and dispatched a technician, who cleaned a strainer that was not very dirty. Nine days later the case came back with one line added: strainer cleaned, fault recurred. The assistant that wrote the summary had taken over first-line support in the autumn of the second year, after the expensive winter. The decision was defensible then and stayed defensible on the dashboard. The assistant resolved about two thirds of contacts itself, answered at night, found the right article among the 4,200 in the knowledge base, and filled in the ticket form more consistently than the agents had. Within a year the first line went from 40 people to 12. What the assistant could not resolve it passed to the second line with a summary of this kind and the full transcript attached below it.

The second time, the case went to Jana. She scrolled past the summary to the transcript, as she usually did. The caretaker had written that the unit locked out “always after the weekend, Monday first thing”. The summary had turned that into three occurrences in ten days, which was correct and no longer said the same thing. Jana had spent years on the first line taking calls from caretakers, and a fault that kept to

the school week did not sound to her like a part. She phoned him and asked what happened in the building on Friday afternoons. The cleaning crew turned the classroom radiator valves down before the weekend, and on Monday morning the system came out of setback against a mostly closed circuit. It took one call and no parts.

The graduate had not been careless. He had compared the summary with the tree, and the two agreed, because both had been drawn from the same knowledge base. Jana had recognized something, and what she recognized was in the caretaker's wording: she had heard enough Monday complaints to know what an ordinary one sounds like and what a failing pump sounds like. For her the summary was one more artifact, an index to the transcript. For him it was the case.

Katrin asked afterwards how many people on the second line would have opened the transcript. The second line stood at 19, with 4 openings that could not be filled from inside. Jana named four: herself, Milan and two other seniors, two of the four due to retire within three years. The 28 jobs that had gone from the first line were the jobs in which she had learned to hear it. By every figure Service kept, response times, article coverage, consistency of categories, it was better informed than it had ever been. The assistant still sat where it had been put a year earlier, between the operator's sentence and the first person at Halvern who read it.

A model changes the task, and where it sits decides which task

Hutchins' method, chapter 7 argued, turns "the organization thinks" into questions that can be asked of any medium in the chain of representations: what does it receive, what does it produce, and can the next medium check it? Applied to a model in the loop, the questions

decide what task the people are left with, and the answer depends on where the model sits.

That the placement is a choice, and whose choice, follows from a sentence of Hutchins' that chapter 7 also used: "Humans create their cognitive powers by creating the environments in which they exercise those powers."¹⁹ A model in the loop is such an environment, made by the organization for its own people, and the organization is responsible for what it does to their task in the same way that it is responsible for a form, a dashboard or a handoff. The difference from those older artifacts is only that the model's effect on the task is larger and less visible, because the work it takes over disappears from view along with the work.

On the input side, at perceiving, a model that surfaces patterns leaves the people with the task of deciding which pattern is a signal and which is an artifact of the model's own categories. That is the partition problem of chapter 2: a model can only find the distinctions it was built to make, and someone chose them. The people downstream need to know what the model cannot see in order to know what its silence means.

In the middle, at understanding, a model that produces an account of the situation, a summary, an analysis, a briefing, leaves the people with the task of checking an account they did not build. This is where Hutchins' concern with inspection bites. On the navigation team the bearing recorder receives a bearing as "strings of spoken digits"²⁰ and can hear whether the number is possible; the plotter can see whether the line crosses the others where it should. Each step of the chain can be checked at the next. An account that arrives finished, from a medium whose working cannot be seen, can be checked only against the reader's own model of the situation. Where the reader has one, the account is a useful draft. Where they do not, it

¹⁹Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

²⁰Hutchins (1995), ch. 4, p. 226.

is the account, and the organization's picture of its situation has been produced by a participant that nobody can question.

On the output side, at choosing, the model proposes. *Punch and Judy* described action selection, by analogy with the vertebrate nervous system, as the problem of “how a field of possible actions is held back, compared, and selectively released”, and noted that in the organism selection works by inhibition: most candidates are held back, and “something has to decide which one is expressed”.²¹ A model that proposes or ranks actions widens the field of candidates and makes them cheap. It does nothing to the release. The book's position, from chapter 6, is that release from inhibition must stay with people who can evaluate what is being released, and the machine has made that condition harder to meet in two ways: there are more candidates, and the people releasing them may no longer be the ones who would have generated them.

One consequence for chapter 6's other term follows. A decision, in the book's sense, is an explicit commitment that others can rely on and later check, and chapter 6 argued that a decision record is what turns an action into a probe: a stated expectation the next round can confirm or refute. When the action was proposed by a model, the record has to hold the expectation of the people who released it, in their own account of the situation, or there is nothing for the outcome to test against except the model's ranking. An organization that records the model's proposal and its own acceptance has recorded the wrong thing. Weick's “Action shapes cognition”²² still holds, but only for whoever held the expectation.

Stafford Beer's redundancy of potential command names the risk in organizational terms. Command, he wrote following McCulloch, “changes from time to time.

²¹*Punch and Judy*, ch. 2 “System Four Is Not a Pair of Eyes”, section “Action selection: the function the diagram omits” (40-stories/projects/punch-and-judy/parts/p01-ch02_system-four-not-eyes.md). The neuroanatomical claims (Grillner, Stephenson-Jones) are used as an analogy.

²²Weick, K. E. (1995). *Sensemaking in Organizations*. Sage, p. 12.

Its location is a function of the information available”.²³ Beer meant this as a strength: the decision migrates to “the little group of employees, possibly very junior, which understands the need”. With a model in the loop the same principle can carry command to wherever the information appears to be, and the information now appears to be in the model. Nothing in the principle says that where the information is, the understanding is too. The Judgment Zones work asks the question in its shortest form, “who decides?”, and defines the agency to do so as “capacity × room × will”.²⁴ A model can be given room. Capacity, in the sense of a tested model of the task, stays where it was grown.

Delegation only as deep as the capacity to evaluate

The limit follows from the internal model, and the book’s own practice states it in one clause: delegation is legitimate “only as deep as the capacity to evaluate what comes back”.²⁵ The clause was written about people delegating to people. It holds without change for delegating to machines, and the reason is the reason the book has been giving since chapter 1.

The depth is set per task, not per organization. A legal team that has drafted a hundred contracts of a kind can delegate the drafting of the next one and read the result as a draft; the same team cannot delegate a kind of contract it has never drafted, because nothing in it would notice what was wrong. A finance function can delegate the forecast it has built by hand for years and still see where the machine’s version departs from what it knows about the business; it cannot delegate a forecast for a business it has never modeled. In each case the question is the same, and it is chapter 12’s question turned round: has anyone in

²³Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

²⁴Judgment Zones client deck v0.4, 06 Judgment Zones/judgment-zones-client-deck-v0.4.md.

²⁵The Tautai company belief core, 01 Tautai/01-Canon/01.Introduction/00.core-belief/00.core-belief.md (technology and authority). Standalone readers: Appendix A.

this role passed far enough along the trajectory to hold a model of the task against which the output can fail?

Evaluating presupposes a model of the task. Beer's sentence that "the regulator is as good as its contained model of the system regulated, and no more"²⁶ was discussed in chapter 1 with the caution the book attaches to it; in the loose sense the book does use, it says that whoever checks the machine's output can check it only as far as their own account of the situation reaches. Beer's other formulation is the more useful one here: "we can only control a system to the extent that we understand it. ... all we're controlling is the projection of the model-in-our-heads."²⁷ An organization that delegates past the depth of its own model is no longer evaluating. It is accepting.

Chapter 9 supplied a test for this that was built for another purpose and fits here without alteration. *Punch and Judy* turned the old complaint that a unit's forward-looking function produces reports nobody acts on into a claim that can be checked: the model attributable to that unit is trivial, because nothing the unit decides depends on it.²⁸ Apply the test to a role that receives model output. Does anything the role decides depend on the role's own account of the situation, or would the same decisions follow from the machine's output whoever sat in the chair? Where the second holds, the role has no attributable model of its own, and the organization has delegated to the depth of zero capacity to evaluate. It may be getting good decisions. It has no way of knowing.

It is sometimes said that the machines have world models of their own, and that these can stand in for the organization's. The phrase can be granted for the sake of argument, and the argument still fails, for two reasons the book has already assembled. Chapter 9 followed Beer in requiring two models of any unit that steers: a model of

²⁶Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 114. On the limits of the warrant see chapter 1.

²⁷Beer (1979), p. 80.

²⁸*Punch and Judy*, ch. 3, section "From a complaint into a claim", on Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025), A "good regulator theorem" for embodied agents, arXiv:2508.06326. (P&J).

its environment and a model of itself. Whatever a model trained outside the organization has learned about the world, it has not learned the second, and the second is where identity, purpose and the question of which state is worth defending live. And chapter 2 followed Ross Ashby in noting that the distinctions a model can make are set by its observer: “the observer and his powers of discrimination may have to be specified”.²⁹ The distinctions in a machine model were set by whoever built and trained it, for purposes that were not the organization’s. To act on that model is to act on someone else’s partition of the world, and not to know it.

Here the two theories of Part IV are needed together, and the chapter keeps them apart as the book has done. For Hutchins, the machine is one more medium in the functional system, and the system, people and artifacts and now models, is what thinks; on that reading a model in the loop is a change in the system’s cognitive properties, to be traced like any other change of medium, and the questions of chapter 7 apply to it. For Nonaka, knowledge “is created by individuals”,³⁰ and judgment, the most tacit form of it, is grown in people along the trajectory chapter 12 described. The book needs both because the machine sits in Hutchins’ system and the capacity to evaluate it sits in Nonaka’s individuals, and an organization has to design for both at once: where the model sits in the chain, and where the people are grown who can check what it hands on.

Hutchins’ sentence therefore still holds, and the chapter can say what the tool has changed the task into. A model in the loop leaves the people with evaluation: of patterns, of accounts, of candidates for action. Evaluation is the work of judgment, and judgment is grown along a

²⁹Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

³⁰Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 17 (knowledge created by individuals; intention, autonomy, fluctuation), p. 18 (the “breakdowns” passage after Winograd and Flores). (public PDF, pp. 14–21).

trajectory the same machine tends to remove. That is the shape of the problem, and it is a design problem. The next chapter takes it up as one: which conversions are left to machines and which are protected as occasions for people; which selections are released by people who can evaluate them; and which roles are kept, by design, as the places where the next people who can evaluate are made.

What would count against it

The chapter claims that machines can take over combination and support externalization, cannot perform socialization or internalization for people, and that an organization which delegates the entry-level tasks without protecting those two modes becomes better informed and less capable, because delegation is legitimate only as deep as the capacity to evaluate what comes back. It would count against this if organizations that delegate combination and interpretation to models without protecting socialization and practice kept their judgment capacity as well as those that protect them; or if judgment grew reliably where the entry roles had been removed. It would also count against it if people without a model of a task turned out to evaluate a machine's output on that task as well as people with one. Where to look: the experience-pipeline argument of *The Tautai Principle* and the longitudinal studies it would need; the expertise literature of chapter 12; and, for the evaluation claim, studies of how experts and novices use the same tool differently.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Cole, M., & Griffin, P. (1980). Cited in Hutchins (1995), ch. 3.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Judgment Zones client deck v0.4.

- Klein, G. (1998). *Sources of Power: How People Make Decisions*. MIT Press.
- Mathis, K. (2025). *The Tautai Principle*.
- McNeese, M. D., Salas, E., & Endsley, M. R. (Eds.) (2021). *Foundations and Theoretical Perspectives of Distributed Team Cognition*. CRC Press.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- *Punch and Judy*.
- Reep, J. A. (2021). Expertise and distributed team cognition: A critical review and research agenda. In McNeese, Salas & Endsley (Eds.), *Foundations and Theoretical Perspectives of Distributed Team Cognition*, ch. 8. CRC Press.
- Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, Communications, and the Public Interest* (pp. 37–72). Johns Hopkins Press.
- Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*.
- The Tautai company belief core.
- Tsoukas, H. (2011). How should we understand tacit knowledge? A phenomenological view. In M. Easterby-Smith & M. A. Lyles (Eds.), *Handbook of Organizational Learning and Knowledge Management* (2nd ed.). Wiley.
- Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326.
- Weick, K. E. (1995). *Sensemaking in Organizations*. Sage.

Chapter 14

Chapter 14 — Designing the understanding organization

A design from 1986

In the January–February 1986 issue of the *Harvard Business Review*, Hirotaka Takeuchi and Ikujiro Nonaka described how a group of companies, observed for the most part in Japan, developed new products. They found six things in common: built-in instability, self-organizing project teams, overlapping development phases, what they called “multilearning”, subtle control, and the organizational transfer of learning.¹ The article is remembered today because Scrum took its name from its rugby metaphor. What interests me here is something else. Eight years before Nonaka set out his theory of knowledge creation, the two authors had described a working design of the loop this book has been taking apart, as a set of decisions companies had made about how their people would notice, understand, choose and learn together.

Two of those decisions are worth holding on to for the rest of this chapter. The first is how learning traveled from one project to the next. It did not travel in documents. The companies transferred it “through ‘osmosis’”, by assigning key people from a finished project to the next one.² The

¹Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116, p. 2 (the six characteristics). That Scrum took its name from the article’s rugby metaphor rests on the article itself — its section heading is “Moving the Scrum Downfield” — and on the six characteristics listed above.

²Takeuchi & Nonaka (1986), reprint p. 8 (subtle control; transfer of learning through “osmosis”).

second is a warning the authors attached to their own findings. “Institutionalization, when carried too far, can create its own danger. Passing down words of wisdom from the past or establishing standard practices based on success stories works well when the external environment is stable.”³ Their companies had therefore made unlearning part of the design. The authors added that their observations were “drawn, for the most part, from observing how the development process was managed in Japan”, and that general conclusions had to be drawn with caution.⁴ I take the caution with the finding.

I begin with this page because it shows what the chapter is for. The thirteen chapters before it have argued that understanding is produced in a loop, that the loop is held by people together with their artifacts, forums and structures, and that what it produces settles as knowledge, skill and, in time, judgment. None of that becomes a capability of an organization until somebody decides where the loop’s parts will sit and who will hold them. Takeuchi and Nonaka’s companies had decided. Most organizations have not; their loop exists, because every organization acts on some account of its situation, and its shape has been settled by accident.

The loop as a design object

Hutchins gave the design argument its foundation in one sentence: “Humans create their cognitive powers by creating the environments in which they exercise those powers.”⁵ If the unit that thinks is the functional system of people and artifacts (chapter 7), then the shape of that system decides what the organization can think, and the shape can be chosen. Beer said the same of the model an organization steers by: it “may be disseminated, in separate chunks of cognition around the firm”, and nobody

³Takeuchi & Nonaka (1986), reprint p. 9 (institutionalization and unlearning; the limits of the sample).

⁴Takeuchi & Nonaka (1986), reprint p. 9 (institutionalization and unlearning; the limits of the sample).

⁵Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, ch. 3, p. 169.

may think of it as a model, “but it must be there”.⁶ Something that must be there, in chunks around the firm, is a design object whether or not anyone has designed it.

Two constraints from the earlier chapters govern how it is designed. The first is recursion (chapter 9). Every unit that faces an environment of its own needs its own understanding, so the design is done per level, for the team, the business unit and the company, and the grids below are filled in for each of them separately. One picture for the whole company hides what its parts need to see. The second is that the design has to work under two readings of where knowledge sits, which the book has kept open since chapter 7. For Nonaka, knowledge is “created by individuals” and the organization amplifies it; for Hutchins, the system of people and artifacts is what thinks, and on the *Palau* it found a solution before any of its members did.^{7,8} A design that only builds the individual’s path, through apprenticeship and practice, leaves the system’s learning unsaved. A design that only builds the system’s representations, forums and records leaves nobody who can grow the judgment the system needs. The grids therefore ask both questions in every row: what the system holds, and who among the people holds it.

Hutchins also removed a division that most organizations build into their org charts. After the *Palau* he concluded that “The processes by which work is accomplished, by which people are transformed from novices into experts, and by which work practices evolve are all the same processes.”⁹ Doing the work, learning to do it and changing how it is done are one process, not three

⁶Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

⁷Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 17 (knowledge created by individuals; intention, autonomy, fluctuation), p. 18 (the “breakdowns” passage after Winograd and Flores). (public PDF, pp. 14–21).

⁸Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

⁹Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

programs owned by operations, human resources and a change office. The six grids that follow are one design for that one process. They are drawn from the earlier chapters, and each is a set of questions to ask of a unit, with the answers written down as they are found before anything is changed.

Grid 1 — Perceiving

Chapter 2 argued that attention is attenuation. Every organization filters, and the design question is whether it filters on purpose. Weak signals are filtered three times: by the model, which cannot register a cue that fits none of its distinctions; by the channels, which decide whether a cue seen at the edge reaches anyone who can act; and by silence, when the people who saw it do not say it. A scanning department addresses none of the three.

Design question	What to look at	Failure it addresses (chapter 2)
What is attended to by design?	Agendas, dashboards, standing forums, escalation paths: which of them were chosen, and which happened	Attention settled by accident; the dashboard encodes an old picture
Which distinctions can the unit make?	The categories in its reports, metrics and forms; what would have no place in them	A cue that fits no distinction is invisible, not weak

Design question	What to look at	Failure it addresses (chapter 2)
Who may raise a cue, and to whom?	The path from the person who sees a change to the person who can act on it; the number of handoffs on it	Signals lost between the edge and the center
How is a cue amplified?	What turns one report into an item somebody has to answer	A signal seen, reported and dropped
Is it safe to report?	What happened to the last person who reported something unwelcome	Silence as the third filter
What do machines watch?	Which patterns a model or dashboard detects, and who reads them	Detection at scale that changes noticing and does not remove it

The row on distinctions is the one organizations skip. Who drew the categories a unit reports in is a decision, usually an old one, and Ashby’s remark that variety depends on “the observer and his powers of discrimination” makes it a design decision.¹⁰ Chapter 10 showed that it is also a question of power. In the grid it is a question of record: who drew the partition, when, and whether anyone since has been allowed to redraw it.

¹⁰Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall, p. 125 (section 7/5), pagination of the 1957 impression as used in *Punch and Judy*, ch. 3.

Grid 2 — Models

The internal model runs through the whole book, and chapters 5 and 9 developed it as orientation and as the core of System Four. The grid for it is the diagnostic that the internal-model note proposed, with one test at its center. *Punch and Judy* turned the old complaint that System Four scans and nothing downstream moves into a claim that can be checked: “the attributed model is trivial”, that is, nothing the unit decides depends on it.¹¹ A unit can hold a strategy paper and a risk register and steer by neither.

Design question	What to look at	Failure it addresses (chapters 5, 9)
Which models does the unit hold?	The accounts it acts on: of its market, its customers, its own capacity	No account at all beyond the plan
Where do they live?	Heads, artifacts (papers, dashboards, spreadsheets), structure (which units exist)	A model in one head that leaves with it
Of what?	Environment and self, separately	An outward model without a model of the unit’s own capacity, or the reverse

¹¹*Punch and Judy*, ch. 3, section “From a complaint into a claim”, on Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025), A “good regulator theorem” for embodied agents, arXiv:2508.06326. (P&J).

Design question	What to look at	Failure it addresses (chapters 5, 9)
Who drew their distinctions?	The people, dates and occasions on which the categories were set	Categories nobody is allowed to change
Does anything decided depend on them?	The last three commitments: which model, if any, they rested on	The attributed model is trivial
How, and how often, are they revised?	The occasions on which the account changed; who may change it	The plan followed as if it were the model

The last row is where Beer’s sentence belongs: “a corporate plan must continuously abort”, or the match between the world and the firm’s notion of it is lost.¹² In the grid it becomes a question about occasions: when did this unit last change its account of its situation, and what made it do so? If the answer is “at the annual planning round”, the unit has a plan and a calendar, and the model is being revised by the calendar rather than by what the unit’s actions produced. Grid 6 takes that up.

At each recursion level the grid is filled in again. The business unit’s model of its market is not the company’s model of the industry, and a design that lets the company’s model stand in for the unit’s has removed the unit’s System Four (chapter 9).

¹²Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

Grid 3 — Choosing

Chapter 6 separated two things that organizations usually call decisions. Action selection is the continuous, mostly tacit choosing that experienced people do inside the loop; Klein’s commanders recognized the situation and the action came with it. A decision is an explicit, accountable commitment that others can rely on and later check. The design question for the first is where it is allowed to happen; for the second, whether it leaves a record.

Design question	What to look at	Failure it addresses (chapter 6)
Where do decision rights sit?	Whether the people who hold the relevant model may act on it, or must escalate	Rights at the top, model at the edge
Is there a commitment forum?	A forum in which competing accounts are compared and then one is committed to, and the two are not the same meeting	Commitment made by whoever spoke first or last
Do decisions carry records?	For each commitment: the expectation, owners, metrics and a review date	Decisions defined in retrospect; nothing to learn against

Design question	What to look at	Failure it addresses (chapter 6)
What releases an action, and what holds it back?	Stop conditions, WIP limits, portfolio reviews; how often “continue watching” is allowed	Zombie projects; everything started, nothing stopped
Where is comparison still required?	Novices, novel situations, justification to others, the complex domain	Recognition trusted where it is confidently wrong
Which selections stay with people?	Where a model proposes or ranks actions, who evaluates and who releases	Release from inhibition handed to a system nobody can evaluate

Beer’s redundancy of potential command is the structural form of the first row. He wrote that “the command centre changes from time to time. Its location is a function of the information available”, and that in a rigid pyramid “there can be no redundancy of potential command”.¹³ Decision rights sit where the model sits, or the organization has built a delay into every decision its edge could have made. The Judgment Zones deck reduces this to a single question, “who decides?”, and to a formula for what a decision right is worth to the person holding it: “Agency = capacity × room × will”.¹⁴

The fourth row draws on a rule from the Tautai cycle review that I return to in the worked example: an experiment may “continue watching” at most once per cycle, because “indefinite watching is how zombie experiments

¹³Beer (1981), p. 232 (McCulloch’s redundancy of potential command).

¹⁴Judgment Zones client deck v0.4, 06 Judgment Zones/judgment-zones-client-deck-v0.4.md.

accumulate”.¹⁵ Whatever the rule, the design needs one, since chapter 6 argued that selection’s default is “no” and something has to release the few actions that are compatible. The record in the third row is what the commitment forum of Sensemaking Article 2 asked for: a “decision record, owners, metrics, review date”.¹⁶ It is the written form of Boyd’s hypothesis, and it is what grid 5 tests.

Grid 4 — Representations and forums

Chapters 7, 8 and 10 gave the design three things that most operating models leave to communication departments: the chain of representations through which a picture travels, the forums in which accounts are compared, and the question of who gives sense. The first is Hutchins’ contribution, and it is what makes the design observable. A bearing on the *Palau* went from a scale to spoken digits to a log to a line on a chart, and each step could be watched and each could fail (chapter 7). In an organization the chain runs from a support ticket to a weekly summary to a slide to a decision, and it can be watched the same way.

		Failure it addresses (chapters 7, 8, 10)
Design question	What to look at	
What is the chain of representations?	Each handoff: from which medium to which, who re-codes, what is dropped	Meaning lost at a handoff nobody owns

¹⁵Tautai cycle review, 01 Tautai/11-Consultant Support/20-Event-Kits/phase-4-intervene-experiment/cycle-review.md.

¹⁶Sensemaking Article 2, 51-Stories/Sensemaking Articles/article-02-sensemaking-moves.md, “Blueprint C — Commitment forum”. Article 3 on decision records as organizational memory.

Design question	What to look at	Failure it addresses (chapters 7, 8, 10)
Which representations are designed?	Problem definitions, decision records, metric dictionaries, boundary objects between units	Representations treated as communication support
How is communication organized?	Who talks to whom, how often, about what; where knowledge overlaps between roles	Dense communication that reinforces one interpretation
Do forums separate frames from commitment?	Whether accounts may compete before one is committed to, and where	Frame competition and commitment in one meeting
Is there a <i>ba</i> for each mode?	A place for shared experience, for articulation, for combining, for practice	Only the systemizing space exists
Who gives sense?	Which social process runs (guided, fragmented, restricted, minimal); who protects reporters	The leader's account as the only one

Design question	What to look at	Failure it addresses (chapters 7, 8, 10)
What is the horizon of observation?	What novices and neighbors can see of each other's work; what tools hide	Learning cut off by the design of a tool

Two of Hutchins' findings decide the third row. Overlapping knowledge is what let the *Palau's* team reconfigure, its members "a sort of flexible tissue" that moved to keep the picture alive; and his simulations of confirmation bias in groups indicated that "more communication is not always in principle better than less".^{17,18} A design that adds channels adds nothing by itself; what matters is whether they carry competing accounts or reinforce one. The fourth row is chapter 8's commitment forum: a forum in which stories are compared, and a separate moment in which one of them is committed to with a record.

The fifth row uses Nonaka, Toyama and Konno's *ba*, the shared context in which knowledge is created, with one type for each conversion mode: originating, dialoguing, systemizing and exercising.¹⁹ Most organizations possess the systemizing space, in the form of a knowledge base and a reporting line, and none of the other three on purpose. The sixth row is chapter 10's: Maitlis' four forms of sensemaking process differ along two dimensions, how much leaders engage in sensegiving (control) and how much stakeholders do (animation), and the design decides which form a unit runs by deciding whether forums stay open and whether the people who report weak signals

¹⁷Hutchins (1995), ch. 4, pp. 221 and 227.

¹⁸Hutchins (1995), ch. 5, pp. 239 and 252.

¹⁹Nonaka, I., Toyama, R., & Konno, N. (2000). SECI, *ba* and leadership: A unified model of dynamic knowledge creation. *Long Range Planning*, 33(1), 5–34. The four types of *ba* (originating, dialoguing, systemizing, exercising) are defined at p. 15 (definitions pp. 15–17), each linked to a conversion mode.

are protected.²⁰ The seventh row is Hutchins' horizon of observation, which "the design of a tool may change" for everyone near it.²¹ A dashboard that shows the result and hides the work has narrowed the horizon for every novice nearby.

Grid 5 — Learning

Chapters 11 and 12 read learning with SECI, as a diagnostic vocabulary: socialization, externalization, combination and internalization, and a spiral from the individual to the group to the organization, "sometimes reaching out to the interorganizational level".²² The diagnosis most organizations receive is the same: strong in combination, weak in the other three. They manage what is already explicit and lose what their people know but have never said. The design grid asks where each mode runs, under what conditions, and along which path a person can grow from novice to judgment.

²⁰Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21–49. p. 21: "four distinct forms: guided, fragmented, restricted, and minimal. These forms result from the degree to which leaders and stakeholders engage in 'sensegiving'"; p. 24: three British symphony orchestras; p. 32 (Fig.): guided → "Unitary, rich account". The two dimensions are control and animation, so "how much leaders and stakeholders control them" was a simplification, corrected in the text.

²¹Hutchins (1995), ch. 6, p. 270.

²²Nonaka (1994), p. 20 (the upward spiral; experimentation and learning by doing).

Design question	What to look at	Failure it addresses (chapters 11, 12)
Which modes run, and where?	Occasions for shared experience (socialization), for putting tacit knowledge into words (externalization), for combining (combination), for practice (internalization)	Combination only
Are the enabling conditions present?	Intention, autonomy and fluctuation at the individual level; redundancy, creative chaos and requisite variety at the organizational level	Learning expected where nothing enables it
Is the trajectory protected?	Entry roles, the simplest parts of the shared task, the horizon of observation, witnessed error correction	Judgment expected without the path that grows it

Design question	What to look at	Failure it addresses (chapters 11, 12)
What is saved in the system?	Whether a solution found under pressure is written into a procedure, a role or a story; whether decision records are reviewed	The <i>Palau's</i> solution, lost
What is unlearned, and when?	Past solutions built into artifacts and procedures; who may retire them	Words of wisdom passed down after the environment moved
Which modes stay with people?	Where machines do combination and support externalization; where socialization and internalization are protected	Better informed and less capable (chapter 13)

Nonaka names intention, autonomy and fluctuation as the conditions at the individual level, and relates fluctuation to breakdowns, which is close to Weick’s interruptions.²³ The organizational conditions, redundancy and requisite variety, and in the research note also creative chaos, are reported from the later pages of the 1994 article; they are in the grid.²⁴ If the reported wording holds,

²³Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37, p. 17 (knowledge created by individuals; intention, autonomy, fluctuation), p. 18 (the “breakdowns” passage after Winograd and Flores). (public PDF, pp. 14–21).

²⁴Nonaka (1994), about pp. 27–29 (organizational enabling conditions: creative chaos, redundancy of information, requisite variety, with reference to Ashby 1956). p. 27: “these enabling conditions—creative chaos,

redundancy is Hutchins' overlapping knowledge as a management principle, and requisite variety is Ashby's law in its design-maxim form. The book uses the law as Ashby wrote it, "variety can destroy variety", and as an argument about whether an arrangement could work, never as a figure.²⁵ The grid asks whether a unit has enough distinct ways of responding to the disturbances it faces; it does not count them.

The third row is where the two readings of the knower meet in practice. Hutchins found that novices "begin by doing the simplest parts of a collaborative-work task", and that "every error-correction event is a learning context not just for the person who commits the error but for all who witness it".^{26,27} The system provides the trajectory, and the individual grows the judgment along it. Chapter 13 argued that AI removes the entry roles at the start of that trajectory and with them the occasions for socialization; the sixth row is where a design protects them on purpose. The fourth row is Takeuchi and Nonaka's osmosis in grid form. What their companies moved to the next project was people, and what the *Palau* failed to move was a solution its people had found and its system had no place for.²⁸

Grid 6 — Rhythms

The temporality thread has run through the book without a home. Sensemaking is retrospective, projection is prospective, orientation is continuous, and the loop has an inner circle that can turn in minutes and an outer one that turns over years (chapter 1). The design variable is

redundancy of information, and requisite variety"; p. 28: redundancy "means the conscious overlapping of company information, business activities, and management responsibilities". Creative chaos is among them; the grid and the outline list it.

²⁵Ashby (1956), p. 207 (section 11/7), pagination as in *Punch and Judy*, ch. 3. "Only variety can absorb variety" is Beer's compression (*Brain of the Firm*, 2nd ed., p. 308).

²⁶Hutchins (1995), ch. 6, p. 264.

²⁷Hutchins (1995), ch. 6, p. 279.

²⁸Hutchins (1995), ch. 8, p. 351 (the solution discovered by the organization and not saved in the system; work, learning and the evolution of practice as one process).

how often each level re-oriens, reviews and unlearns, and who sets that rhythm.

Design question	What to look at	Failure it addresses (thread)
How often does each level re-orient?	The occasions on which a unit revisits its account of its situation, and whether the environment or the calendar sets them	The model revised once a year
When are decisions reviewed?	Whether the review date in each record is honored, and what happens at it	Records written and never read
How often is the portfolio pruned?	The cycle in which experiments are continued, stopped or watched once more	Indefinite watching
When is something unlearned?	A moment, at any level, at which a past solution may be retired	Institutionalization carried too far
Do the levels' rhythms fit each other?	Whether a unit that re-oriens weekly reports into one that re-oriens yearly	The fast level waiting on the slow one

Design question	What to look at	Failure it addresses (thread)
Does the outer circle turn at all?	The path from what was learned to changed capability and changed attention	Faster at acting on a picture nobody improves

Two of the earlier chapters set the limits. Beer’s aborting plan says the rhythm of re-orientation is set by the mismatch between the world and the firm’s notion of it, not by the planning calendar.²⁹ Nonaka treats double-loop learning as “a daily activity for the organization”, which in internal-model terms means that the model is revised in the course of work and not in an offsite.³⁰ The design’s job is to make the occasions exist. A unit that has a commitment forum with review dates, a cycle review with a rule for watching, and a moment for retiring old solutions has a rhythm. A unit that has an annual plan has a calendar.

Measuring understanding without the measurement trap

Chapter 1 described the measurement trap: a metric system encodes the picture that was current when it was designed, and by the time a change shows up in it the change is old.³¹ Designing the loop invites a second trap, which is worse because it looks like rigor. The loop has properties that can be counted, above all how fast it turns, and once they are counted they become targets.

I have set such targets myself. *The Tautai Principle* treats the observe–orient–decide–act loop as the primary answer

²⁹Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley, p. 183. Quoted in full in chapter 1.

³⁰Nonaka (1994), p. 19 (four modes; double-loop learning as a daily activity).

³¹Mathis, K. (2025). *The Tautai Principle*, ch. 4, p. 72. Paraphrased in chapter 1.

to a gap in speed, proposes sensemaking-velocity indicators, sets a decision-latency benchmark of “< 72 hrs” and calls decision latency a “direct proxy for judgment efficiency”.^{32,33,34} Chapter 6 gave the reason this does not hold. Klein’s commanders were fast because recognition ran on a good model, and nobody had pressed them for speed. Latency is an outcome of model quality. Set as a target, it invites fast decisions on poor models, and it rewards the unit that stops asking whether the picture fits. The company belief core that this book grew out of says why the target was wrong in the first place: “much of the work worth doing is not in a race”.³⁵ Appendix A records the three passages as one edition trigger.

The alternative follows from Hutchins. If cognition is the propagation of representations across media, it can be observed, and observation gives measures of a different kind: traces rather than rates. The useful questions are about the loop’s conditions, and each grid supplies one.

- **Do signals travel?** Take a cue that somebody at the edge of the unit saw in the last quarter and follow it. Count the handoffs, name the place where it was dropped, and ask whether the person who saw it would report the next one. That is grid 1, measured by walking the chain.
- **Do decisions carry records?** Of the unit’s commitments in the period, how many have a record with an expectation, owners, metrics and a review date, and how many of the review dates were kept. That is grid 3 and grid 6 together.
- **Do models get revised?** When did the unit last change its account of its situation, what made it do so,

³²Mathis (2025), p. 78 n. 40 (the OODA loop as the primary solution to a speed gap) and p. 80 (sensemaking-velocity indicators) (ch. 4 of the printed book); the two passages are paraphrased here.

³³Mathis (2025), p. 69.

³⁴Mathis (2025), p. 201.

³⁵The Tautai company belief core (“probes with stop conditions rather than plans with milestones”; “much of the work worth doing is not in a race”; delegation only as deep as the capacity to evaluate what comes back).

and did any commitment since then depend on the change. A unit whose answer is “never, and nothing” has the trivial model of grid 2.

- **Do the four modes run?** Where in the unit’s week does shared experience happen, where does someone put tacit knowledge into words, and are the entry roles still there. That is grid 5.
- **Has learning happened?** Beer gave a criterion that needs no number: when the models that a pupil and a teacher hold of each other converge, “learning has occurred”.³⁶ For a unit and its environment the question is whether the unit’s model and what its actions produce have been converging or drifting apart.

Speed and latency stay in the picture, as outcomes to watch. If decisions in a unit became faster after its commitment forum and its records were installed, that is worth knowing, because it is what the argument predicts. If they became faster after a latency target was set, it says nothing about understanding.

One further figure is refused. Requisite variety tempts designers to count states on both sides of a boundary. *Punch and Judy* records the one careful attempt to operationalize the law, Pfiffner’s, and notes that it ended in a judgment of adequacy rather than a count; the law is a priori and licenses an argument, not a number.³⁷

None of these measures is a validated instrument. They are the questions the book’s argument leads to.

How to start: a 90-day shape

The grids describe an end state, and no organization installs six grids at once. The shape below is mine, built from the grids; the Sensemaking Articles carry forum blueprints

³⁶Beer, S. (1979). *The Heart of Enterprise*. Wiley, p. 235.

³⁷*Punch and Judy*, ch. 3, section “What the law becomes when someone has to measure it”, on Pfiffner (2017), p. 165; the a-priori character of the law, Ashby (1956), p. 208 (section 11/10). (P&J)

and shapes of their own that I have not reproduced here.³⁸ It has three moves, each about a month long, for one unit at one recursion level, chosen because it faces an environment of its own and its people will talk.

The first month is a trace, and nothing is installed.

The six grids are filled in for the unit as it is found. Somebody follows a few cues from the edge to wherever they stopped, reads the last commitments and asks what record and what model each rested on, and walks the chain of representations from the first ticket to the last slide. The result is a description of the unit's loop as it exists, which nobody has seen whole before. The month ends with the unit's own people correcting it, because the description is itself a representation and the people who work the chain know where it is wrong.

The second month installs the two cheapest elements.

The first is a commitment forum with decision records: a place where accounts may compete, a separate moment at which one is committed to, and a record with an expectation, owners, metrics and a review date. The second is a channel for cues with a named receiver, with the person who raises a cue told what happened to it and protected if it was unwelcome. Neither needs a reorganization. Both change what the unit notices and what it can later learn from, and both produce the traces the measures above need.

The third month holds the first reviews. The first review dates fall due, and the records are read against what happened. The first cycle review is held with a rule for watching, so that something is stopped. One past solution is chosen for retirement, which is harder than starting anything and is the part organizations skip. And the unit decides, in writing, which selections and which conversions stay with people: what a model may propose, who evaluates it, and where the entry roles for the next novice are. The month ends with the grids filled in a second time.

³⁸Sensemaking Articles 2 and 3 (51-Stories/Sensemaking Articles/) for the commitment forum and decision records; the forum blueprints and 90-day shapes of Articles 2-4. The shape given here is built from the grids.

After ninety days the unit has a loop with a rhythm, a record, and a trace of what changed, at one level. The next move is the next level of recursion, and it is the same three months for the unit above or below, with the interface between them designed on purpose. What travels from the first unit to the second is what traveled in Takeuchi and Nonaka's companies: people who have run the loop once, by osmosis, and not a manual.

The Tautai operating model as one worked example

The operating model I described in *The Tautai Principle* in 2025 is one way of building the six grids into an organization. It was built before this book, from practice rather than from the theories the book has assembled, and it uses names of its own. Appendix A maps those names to the chapters; here I read the model as a worked example of the design, and say where the book has changed it.³⁹

The **Scanning System** and its four foundations are grid 1 and part of grid 4. They answer the chapters on perceiving, situation awareness and sensemaking, and the chapter on groups (chapters 2, 3, 4, 8): what a unit looks at, who may raise a cue, how a picture is held and how it is repaired when it breaks. Read with this book, the system's job is to keep the three filters of chapter 2 open, and its link to chapter 8 is where the commitment forum sits. The printed book gives Endsley's three levels without attribution on its page 74; this book attributes them, and Appendix A records the correction.

The coupled cycles the printed book calls **GOOD** and **FAEE**, one for strategy and one for leadership, are grids 2, 3, 5 and 6 together, and they answer chapters 5, 6 and 11. The GOOD cycle puts a step before observing, the Game, in

³⁹Mapping of the Tautai terms to the chapters as given in Appendix A: Scanning System and its four foundations ▫ chapters 2, 3, 4, 8; OODA and GOOD/FAEE ▫ chapters 5, 6, 11; Superior Sensing, Distributed Response, Intelligence Networks ▫ Part III; Leadership System ▫ chapter 10; capabilities ▫ chapter 12; Steering with Stories ▫ chapters 4, 8. Mathis (2025), ch. 4, 12, 13, 16 as listed in the chapter abstract; the names of the four foundations are not given in the outline and are not given here.

which identity and purpose are set; chapter 10 placed that step where the book's argument needs it, since a model of the environment is useless without a model of the self that is to act in it. The cycle review inside these loops carries the rule quoted in grid 3: an experiment may "continue watching" at most once per cycle.⁴⁰ In the terms of chapter 6 that is a release-and-inhibition mechanism: stopping something is a decision with a record, and starting something rarely is. The printed book also calls the coupling of the two cycles a double loop; Argyris and Schön's double-loop learning is a different thing, and chapter 11 kept the two apart. Appendix A records that as the third departure.

Superior Sensing, Distributed Response and Intelligence Networks are Part III in the model's vocabulary, and they fill grid 4: where understanding is held, how communication is organized, and how the levels of recursion talk to each other. The printed book gathers them under the name of the organization's Nervous System; that grouping is not in the outline of this book.⁴¹ The **Leadership System** is chapter 10 and the last rows of grid 4: who gives sense, which forums stay open, who protects the people who report. Nonaka's middle-up-down management, in which middle managers link the top's vision to the front-line's knowledge, would give the Leadership System its clearest theoretical warrant.⁴² **Steering with Stories** is the narrative thread of chapters 4 and 8: accounts travel as stories, and Nonaka's externalization begins with metaphor,

⁴⁰Tautai cycle review, 01 Tautai/11-Consultant Support/20-Event-Kits/phase-4-intervene-experiment/cycle-review.md.

⁴¹The grouping of Superior Sensing, Distributed Response and Intelligence Networks under the Nervous System is not in the outline of this book, which names the Nervous System for this chapter without mapping it.

⁴²Nonaka (1994), about pp. 29–30 (middle-up-down management, citing Nonaka 1988b). p. 30: "top management provides 'visions for direction' ... Middle management translates these visions into middle-range visions ... while top management articulates the dreams of the firm, lower managers look at the reality. The gap between these two forms of perspectives is narrowed by and through middle management."

which is why a story carries what a dashboard loses.⁴³ The **capabilities**, from knowledge through skill, proficiency and fluency to judgment, are grid 5 and chapter 12, and the experience pipeline argument of the printed book, that removing entry roles removes the path along which judgment is grown, is the row on protected trajectories.⁴⁴

The model therefore covers the six grids. Where this book departs from it is in what the model was asked to deliver. The printed book asked for speed, in the three passages Appendix A lists. This book asks for a valid model, kept alive, and reads speed as what a valid model produces. The change is small in the text and large in practice: a unit that installs the Scanning System to become faster will measure latency and get it; a unit that installs it to keep its picture current will measure whether signals travel and whether models get revised. The diagnostic that goes with the model, which sorts an organization's orientation between navigation and planning, asks the second kind of question and not the first.⁴⁵

Episode — The redesign, grid by grid

In the spring of the fourth year the first product of the redesign at Halvern was an eleven-page description of After-Sales Service as it stood, and for a month it was the only one. Two people from service management followed three cues from the edge of the unit to the place where each had stopped, read the unit's last three commitments, and walked the path from an operator's call to the line on the board slide. Then the pages went to the people who worked that path. Tobias corrected the first of them. The trace had the short-cycling cue stopping at his team lead, and he

⁴³Nonaka (1994), pp. 20–21 (metaphor as the first step of externalization). *The Tautai Principle*, ch. 12 “Steering with Stories” as listed in the abstract.

⁴⁴Mathis (2025), pp. 199–200 (the experience pipeline). The five capabilities follow the Tautai capabilities drawing “Learning — SDL → Capabilities” and its fluency levels, as listed in chapter 12.

⁴⁵The skill *tautai-diagnostic-part1* (assessment of an organization's orientation between navigation and planning).

said it had stopped earlier, at the form, when he picked installation error. Jana corrected another. The trace began the chain at the escalation summary written by the assistant; the operator's own words came before it, she said, and the second line no longer saw them. The second month installed two things: a channel for cues with a named receiver, and a split in the agenda of the Product Quality Board, with a record for each commitment. In the third month the first review dates fell due. By the end of the ninety days the unit differed from the traced one in the following ways.

Grid 1, perceiving. The ticket form kept its five cause categories and gained a sixth, *interaction between systems*, and Henrik's department added the matching root-cause code. Beside the categories stood a box, *does not fit*. A ticket marked that way went with its free text to a second-line specialist on a weekly rota, who had to tell the sender within five working days what had become of it. The form now named who had drawn its categories and who could redraw them: Service and Quality together, once a quarter.

Grid 2, models. The trace had asked which account each of the last three commitments rested on, the Academy, the assistant and the reduction of the first line, and had found none written down. Three accounts were now written, two pages each, with an owner and a date. The account of the installed base recorded what each C4 unit was connected to. The account of why service contracts were not renewed drew on a line of free text that Sales had agreed to add under the CRM's lost-reason dropdown. The account of the unit's own capacity counted the second-line specialists who could take a case that left the decision trees: four of 19, two of them due to retire within three years.

Grid 3, choosing. Provisional service bulletins had carried a record since the second year. Nine had been issued, the Board had confirmed all nine, and none

had been read afterwards against what happened on the units. Each now carried a review date and an expectation that could be checked on that date. An open cue could be kept “watching” for one monthly cycle; after that it was acted on or closed with a reason.

Grid 4, representations and forums. Before each Board meeting the second line and Quality now wrote their root-cause statements apart, and the meeting began by comparing them. Commitment became a separate item at the end, with a record. An account of a failure pattern came to the Board with at least two traced cases and with the person who had seen them. On the second line’s screen the operator’s own words appeared beside the assistant’s summary.

Grid 5, learning. Ride-alongs returned: every technician spent two days a quarter on another technician’s jobs. An escalation that left the trees ended with twenty minutes in which the senior told the Academy graduate what had been asked and why. The first line was rebuilt as an entry role. Four trainee posts were added to the remaining 12, and the trainees read the assistant’s summaries against the calls. Each graduate and trainee attended ten partner commissionings a year beside a senior technician. The knowledge base stopped counting an opened bulletin as one taken up, which was the past solution the unit chose to retire.

Grid 6, rhythms. The receiver read cues and free text weekly. The Board compared accounts, held the review dates and cleared the watching list monthly. The three written accounts and the categories were to be revisited each quarter.

The following winter gave one instance of each kind. In January a technician in region North marked a ticket *does not fit*: an outdoor unit iced on one side where an acoustic enclosure had been fitted. The receiver found four more such tickets in that week’s free text and told her so on the third day. A provisional bulletin with its conditions and a review date went out

nine days after her ticket. In February service management read the quarter's lost-reason lines. Of 41 contracts not renewed, 26 carried a sentence about uptime guarantees or remote monitoring, and the written account of non-renewal was changed from price to that. The unit committed, with a record, to offer a monitored contract at 30 sites and to review it in June. In March one of the first-line agents, who had spent the winter reading summaries against calls and sitting in debriefs, moved to the second line. It was the first of the four openings to be filled from inside. Other parts did not hold. The quarterly review of the three accounts due in January was not held, because the heating season took the people who would have sat in it, and the capacity account was nine months old when the unit committed specialists to the monitored contract. The ride-alongs cost about 1,400 technician-days a year. In the autumn budget round Finance asked what four trainees were for when the assistant resolved two thirds of contacts, and two of the four posts were kept. Nor do the figures settle much. Nine days cannot be set against the eleven weeks of the short-cycling fault, since the fault was a different one and the people handling it knew the channel was being watched. Whether the 26 lines will have changed any renewal is not known before June. This was one unit over one winter, and no comparable unit was left unchanged beside it.

What would count against it

The chapter claims that building the loop into an operating model, per recursion level and in the six grids, changes what an organization notices, chooses and learns, and that measuring the loop's conditions serves adaptation better than setting speed targets. Two things would count against that. The first is the design elements being installed, the forums, records, channels, trajectories and rhythms, without any change in what the organization

notices, chooses and learns; the traces described above would show it, because the cues would still stop where they stopped and the records would go unread. The second is organizations that measure the loop's conditions doing no better, in comparable complex settings, than organizations that set speed and latency targets. Where to look: the high-reliability literature and the accident analyses named in chapter 1, and the follow-up of organizations that adopted the design, including the spine case once it is written.

The last of Takeuchi and Nonaka's six characteristics was the organizational transfer of learning, and their last warning was about institutionalization. A design of the loop is subject to its own argument. The grids are an account of how understanding is produced, and an account is a model that has to keep aborting. Appendix A records where this book's account has already departed from the one I printed in 2025; the next edition will record where it departs from this one.

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21–49.
- Mathis, K. (2025). *The Tautai Principle*.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Nonaka, I., Toyama, R., & Konno, N. (2000). SECI, *ba* and leadership: A unified model of dynamic knowledge creation. *Long Range Planning*, 33(1), 5–34.
- Pfiffner (2017).

- Takeuchi, H., & Nonaka, I. (1986). The new new product development game. *Harvard Business Review*, January–February 1986, reprint 86116.
- Virgo, N., Biehl, M., Baltieri, M., & Capucci, M. (2025). A “good regulator theorem” for embodied agents. arXiv:2508.06326.

Chapter 15

Appendix A — This book and The Tautai Principle

This book is standalone. It does not assume *The Tautai Principle*, and its chapters use general terms for the things the two books share. This appendix serves a reader who knows both. It maps this book's terms to the Tautai operating model's, so that an argument can be carried from one book to the other; and it consolidates the edition triggers — the points where this book departs from the printed book, and where a future edition of *The Tautai Principle* will have to decide.

Nothing here overrides the package canon. The canon and the printed book govern the Tautai terms; where this book departs from the printed book, the departure is recorded as an edition trigger and kept visible rather than merged away. Chapters 6 and 14 point here for the first and largest of them.

Terms mapped between the two books

The table reads left to right: a term this book uses, and the Tautai term for the same thing. Where the two differ, the note in the right-hand column says how.

This book	<i>The Tautai Principle</i> / package canon
Understanding — the umbrella term for the whole; sensemaking reserved for Weick’s activity	The Scanning System; <i>sensemaking</i> used for the whole activity. <i>Differs:</i> the book uses <i>understanding</i> for the whole and reserves <i>sensemaking</i> for what people do when the picture breaks (chapter 4).
The internal model	System Four — the model the organization steers by. <i>Differs:</i> the book develops the internal model as the core of System Four and keeps its loose and strong senses apart (chapters 5, 9).
Action selection (continuous, tacit) and decision (explicit, accountable commitment)	“Decide” in the GOOD cycle. <i>Differs:</i> the book splits choosing into two terms; the GOOD cycle carries one Decide step, and the VSM diagram omits action selection (chapter 6).
Judgment	Capabilities — the top of the ladder. <i>Note:</i> the book reads judgment as a rich, practice-tested, largely tacit internal model (chapter 12).

This book	<i>The Tautai Principle / package canon</i>
The skill-development cycle (SDLC)	The “ Learning — SDLC → Capabilities ” drawing. <i>Note:</i> the name is shared — here “SDLC” is the skill development cycle. The five steps and their order are settled — knowledge, skill, proficiency, fluency, judgment — and the cycle is read as a design; the cycle’s own order and name are settled (chapter 12).
Requisite variety	Requisite variety. <i>Note:</i> no divergence — both use it as a design argument about what could work, never a figure (chapters 2, 9).
System Four	System Four; Intelligence Networks. <i>Note:</i> the book treats System Four as modeling and anticipation, not “a pair of eyes” (chapter 9).
Weak signal detection	Weak Signal Detection — one of the four foundations of the Scanning System. <i>Note:</i> the same name; the book treats detection as a social activity, and does not use the draft’s “Weak Signal Sensing” (chapter 2).

This book	<i>The Tautai Principle</i> / package canon
The four sisters of understanding	Weak Signal Detection · Situational Awareness · Collective Sensemaking · Attention Management. <i>Differs:</i> the four are covered here as chapter activities (chapters 2, 3, 4, 8), not gathered under one name; the names used here are the book’s own.

The structural mappings follow the same pattern. The book’s perceiving, understanding and collective sensegiving (chapters 2, 3, 4, 8) are the **Scanning System and its four foundations**; choosing and committing (chapter 6) is the **GOOD cycle** (Game, Observe, Orient, Decide); acting and reviewing is the **FAEE cycle** (Focus, Align, Enable, Empower); the distributed forms of understanding (Part III) — gathered in the printed book under the organization’s **Nervous System** — are **Superior Sensing, Distributed Response** and **Intelligence Networks**; leadership (chapter 10) is the **Leadership System**; narrative (chapters 4 and 8) is **Steering with Stories**; and learning and judgment (chapters 11 and 12) are the **Learn arc and capabilities**. Chapter 14 shows the whole design in the Tautai frame as one worked example.

Where this book departs from the printed book

The project records three departures, kept as one consolidated list of **edition triggers**. Each says what the printed book does, why this book departs, and what the correction is.

1. Speed and latency as targets. In three places *The Tautai Principle* treats speed or latency as the thing to be aimed at: OODA as “the primary solution to the Speed Gap” (p. 78 n. 40); “Sensemaking Velocity” indicators (p. 80);

and a decision-latency benchmark of “< 72 hrs” (p. 69) with latency called a “direct proxy for judgment efficiency” (p. 201). This book departs because it reads speed as a consequence of the quality of the model, not as the premise of the loop (chapter 6; the epistemic premise; the core belief that much of the work worth doing is not in a race). The correction is a change of practice rather than a new figure: watch speed, do not set it. The measures worth keeping are the loop’s conditions — whether signals travel, whether decisions carry records, whether models are revised — and latency is an outcome to observe, not a target to set (chapter 14).

2. Endsley’s three levels without attribution. The printed book states situation awareness in Endsley’s three levels — perception, comprehension, projection — without attributing them to her (p. 74). This book departs because it fixes the vocabulary to its source: chapter 3 names Endsley and separates the three levels from the book’s own reading of where awareness fails.

3. “Double loop” with two meanings. The printed book uses “double loop” for the coupled GOOD/FAEE cycle, while the literature uses the same words for Argyris and Schön’s double-loop learning. This book departs because the two are different things under one name: it names both and keeps them apart (chapter 11).

Draws on

The Tautai Principle (Mathis, 2025; 61 Books/_Published/The-Tautai-Principle.pdf) for the page locations as listed in the outline and the evidence records; the edition triggers in the project’s PLAN.md (row R); the outline; and the authority map, which governs the two works.

References

- Mathis, K. (2025). *The Tautai Principle*.

Chapter 16

Appendix B — Glossary

This is the book's glossary: the terms the chapters define, or fix in a particular sense, each given a short definition in the book's own vocabulary and the chapters where it is developed. A definition states the sense the book uses, not the term's whole history in the literature. For the terms this book shares with *The Tautai Principle*, and for the points where it departs from the printed book, see Appendix A.

Terms

- **Action selection** — The continuous, mostly tacit choosing that goes on inside the loop, as distinct from a decision. Among the available actions most are incompatible, and something has to release one; the book reads selection as release from inhibition, whose default is “no” (chapter 6).
- **Artifacts** — The tools, forms, templates and records an organization thinks with. Artifacts hold the solutions of earlier generations, which is both capacity and inertia (chapter 7).
- **Attention** — Attending and noticing; the input-side selection problem. Attention is limited in people and in organizations, and in cybernetic terms it is attenuation — every organization filters variety, and the question is whether it filters deliberately (chapters 1, 2).

- *Ba* — Nonaka, Toyama and Konno's term for the shared context in which knowledge creation happens; the book uses it for the context a forum provides (chapters 8, 11).
- **Commitment forum** — A forum that separates the competition between frames from the commitment to one of them: a place where accounts may compete, and a separate moment at which one is committed to (chapters 6, 8, 10).
- **Cynefin** — Dave Snowden's sensemaking framework, used here for the complex domain, where probing is the appropriate mode; that recognition there can be confidently wrong is the book's own reading of the framework, not a statement of Snowden and Boone's (chapters 1, 6). Snowden & Boone (2007).
- **Decision** — The explicit, accountable commitment that others can rely on and later check, as distinct from action selection. Without a record, decisions get defined in retrospect (chapter 6).
- **Decision record** — A record of a decision carrying an expectation, owners, metrics and a review date. It turns an action into a probe, and it is what learning tests against (chapters 6, 11).
- **Distributed cognition** — Hutchins' approach: the unit that thinks is the functional system of people and artifacts, and cognition is the propagation of representational state across media, which makes it observable (chapter 7).
- **Double-loop learning** — Argyris and Schön's distinction: single-loop learning corrects the action within the organization's existing governing assumptions, while double-loop learning questions and changes the assumptions themselves. In the book's vocabulary, single-loop revises the action and double-loop revises the model (chapter 11).
- **Enactment** — Weick's: organizations do not simply meet an environment that is out there; they help produce the environment they then try to read (chapters 1, 4).

- **Equivocality** — The condition in which the cues support several plausible readings at once, so that more information does not help because nothing is missing. Weick distinguishes it from uncertainty, where information is missing and can be gathered (chapter 4).
- **The four sisters of understanding** — The four activities named across the canon — Attention Management, Situational Awareness, Collective Sense-making and Weak Signal Detection. The book covers them as chapter activities rather than under one name; the names used here are the book's own (chapters 2–4, 8; Appendix A).
- **Functional system** — People, artifacts, routines and roles together, taken as the unit that thinks (chapter 7).
- **Internal model** — The account an organization acts on, of its environment and of itself, held in heads, artifacts and structure. In the loose sense every organization acts on some account; in the strong sense a regulator needs a model of what it regulates — the first is uncontroversial, the second is argued, not proved (chapters 1, 5, 9).
- **Judgment** — Knowing what the situation calls for when the procedure does not say, including when the procedure should be set aside and when nothing should be done. A rich, practice-tested, largely tacit internal model, and the most tacit and hardest to hand on of the five capabilities (chapter 12). Polanyi (1966); Dreyfus & Dreyfus (1980, 1986).
- **Learning loop** — Nonaka's spiral: the movement of knowledge through the four conversion modes and up the levels, a loop because internalized knowledge becomes the tacit base for the next round of socialization (chapter 11). See also double-loop learning and the loop.
- **The loop** — The book's central figure: perceiving, understanding, choosing, acting and learning drawn as one loop, with a fast inner circle from cue to account

to action and a slow outer circle through learning and capability (chapter 1).

- **Orientation** — Boyd's function at the center of his loop: the continuous destruction and creation of concepts to match observed phenomena with mental models — that is, the continuous revision of the internal model (chapter 5).
- **Partition problem** — The cybernetic point that the distinctions a model can make are set by an observer, and someone chose them. What can be attended to is bounded by the distinctions the model can make, and who draws them is a decision (chapters 2, 10; Ashby 1956).
- **Prospective sensemaking** — The forward-looking counterpart of sensemaking: the construction, in conversation, of images and schemes of future objects and phenomena (chapter 4).
- **Recognition-primed decision** — Klein's finding that experienced people recognize the situation as a type and a workable action comes with it, so that comparing options is rare. The basis of the book's point that understanding already contains the choice (chapter 6).
- **Recursion** — The principle, from the Viable System Model, that every unit which faces its own environment needs its own understanding, and so its own loop (chapter 9).
- **Redundancy of potential command** — Beer's principle, after McCulloch, that command changes from time to time and its location is a function of the information available; the structural form of the rule that decision rights sit where the model sits (chapters 6, 10).
- **Representation** — The medium a computation is carried in. A representation is re-coded for a new medium at each handoff, and meaning is translated, lost or corrected on the way (chapter 7).
- **Requisite variety** — Ashby's principle, used in the book as an a priori design argument about which arrangement could work, never as a figure; the book

reports no variety numbers (chapters 2, 9; Ashby 1956).

- **SECI** — Nonaka and Takeuchi's theory of organizational knowledge creation: knowledge exists in tacit and explicit form, and the four modes of conversion — socialization, externalization, combination, internalization — move it between forms and up the levels in a spiral. The book uses the modes as a diagnostic vocabulary, not as a claim that tacit knowledge is literally converted (chapters 11–14). Nonaka (1994).
- **Sensegiving** — Gioia and Chittipeddi's term for the leader's side of sensemaking: leaders give sense as much as they make it. The social processes of sensemaking differ by how much leaders and stakeholders control them (guided, fragmented, restricted, minimal) (chapter 10). Gioia & Chittipeddi (1991); Maitlis (2005).
- **Sensemaking** — Reserved in this book for Weick's activity: what people do when the situation fits no picture they have, in equivocality, surprise and interruption. A social, retrospective construction of a plausible account good enough to act on, which action then tests; not analysis, and not the interpretation of given data (chapter 4).
- **Situation awareness** — Endsley's term for holding a picture that fits, with three levels: perception of the relevant elements, comprehension of what they mean together, and projection of their state in the near future (chapter 3). Endsley (1995); Endsley & Jones (2012).
- **The skill-development cycle (SDLC)** — Reading the five capabilities — knowledge, skill, proficiency, fluency and judgment — as a cycle from what can be taught to what can only be grown. The order of the five is settled (proficiency precedes fluency; judgment comes after fluency) and the cycle is read as a design; the cycle's own order and name are settled (chapter 12).
- **System Four** — In Beer's Viable System Model, the organization's outward modeling and anticipating

function: it must contain a model of the organization and of its environment. It is not “a pair of eyes”, and its testable failure is a trivial model, one on which nothing decided depends (chapter 9; Beer 1981, 1979).

- **Tacit and explicit knowledge** — Polanyi’s distinction, following which Nonaka defines explicit knowledge as transmittable in formal, systematic language, and tacit knowledge as personal, hard to formalize and communicate, and deeply rooted in action, commitment and involvement in a specific context (chapters 11, 12).
- **Transducer** — Beer’s term for what converts a representation from the language of one system into that of another at a boundary. Hutchins supplies what happens inside one: a state in one medium becomes a state in another, and something is preserved and something dropped (chapters 7, 9).
- **Understanding** — The book’s umbrella term: how an organization comes to know what situation it is in and what to do, and how that knowing is kept alive. Where the text needs a word for the whole, it says understanding; *sensemaking* is reserved for Weick’s activity and *situation awareness* for Endsley’s (chapter 1).
- **Unlearning** — The deliberate retirement of a past solution or governing assumption; the part organizations skip, made harder because past solutions are built into artifacts and procedures (chapters 11, 14).
- **Viable system / Viable System Model** — Beer’s model of the functions any viable system needs (Systems One to Five), and of the boundaries, channels and transducers between them; the source of the book’s vocabulary for structure as cognitive architecture (chapter 9).
- **Weak signal** — An early, ambiguous cue of change. Weak signals are not rare but filtered — by the model (a cue that fits no distinction is invisible), by channels, and by silence (people who see them do not say

them) — so detecting them is a social activity, not a scanning department (chapter 2).

References

- Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.
- Beer, S. (1979). *The Heart of Enterprise*. Wiley.
- Beer, S. (1981). *Brain of the Firm* (2nd ed.). Wiley.
- Dreyfus, S. E., & Dreyfus, H. L. (1980). *A Five-Stage Model of the Mental Activities Involved in Directed Skill Acquisition*. Operations Research Center, University of California, Berkeley.
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over Machine*. Free Press.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64.
- Endsley, M. R., & Jones, D. G. (2012). *Designing for Situation Awareness* (2nd ed.). CRC Press.
- Gioia, D. A., & Chittipeddi, K. (1991). Sensemaking and sensegiving in strategic change initiation. *Strategic Management Journal*, 12(6), 433–448.
- Maitlis, S. (2005). The social processes of organizational sensemaking. *Academy of Management Journal*, 48(1), 21–49.
- Nonaka, I. (1994). A dynamic theory of organizational knowledge creation. *Organization Science*, 5(1), 14–37.
- Polanyi, M. (1966). *The Tacit Dimension*.
- Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76.