

ECSTATIC ECONOMICS · THE LUMINOUS BUSINESS METHODOLOGY

The Strategy That Reaches the Floor

Where a Plan Is Lost Between the Board and Monday Morning

Ammanuel Santa Anna

Santa Anna, A. (2026). The Strategy That Reaches the Floor: Where a Plan Is Lost Between the Board and Monday Morning. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989179
Luminous Library Licence 1.0 · business.hautelumiere.life/library/the-strategy-that-reaches-the-floor

THE LICENCE

The Strategy That Reaches the Floor
Where a Plan Is Lost Between the Board and Monday Morning

Ecstatic Economics · The Luminous Business Methodology. The Haute Lumière edition, 2026. © 2026 Ammanuel Santa Anna.

Released under the Luminous Library Licence 1.0 (hautelumiere.life/library-licence): free to read, share, quote, translate and teach, non-commercially, for ever. The Author keeps every other right, including every commercial right.

Published by Luminous Prosperity Inc. for Haute Lumière. Author ORCID 0009-0008-3244-0306.

DOI 10.5281/zenodo.23989179 · this record's page: business.hautelumiere.life/library/the-strategy-that-reaches-the-floor/

Please cite as: Santa Anna, A. (2026). The Strategy That Reaches the Floor: Where a Plan Is Lost Between the Board and Monday Morning. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989179

A NOTE BEFORE YOU READ

This work is offered for education and leadership development. It is not legal, financial, accounting, investment, medical or psychological advice; take the advice of a qualified professional who knows your circumstances. Companies and events are described from the public record as the author understands it; characterisations are the author's analysis. Examples are illustrative, and no particular outcome is promised to any reader.

ABSTRACT

Where a plan is lost between the board and Monday morning: the card, the middle, the target, the review and the pattern. Where a plan is lost between the board and Monday morning: the card, the middle, the target, the review and the pattern; seven chapters with a field manual.

CONTENTS

CHAPTER ONE Three Priorities on a Card 4

CHAPTER TWO The Ball Thrown Back 16

CHAPTER THREE The Number That Stood In 28

CHAPTER FOUR The Review That Reported Green 39

CHAPTER FIVE The Strategy You Already Follow 50

CHAPTER SIX The First Ninety Days 63

CHAPTER SEVEN The Field Manual 76

NOTES References 90

CHAPTER ONE

Three Priorities on a Card

A strategy reaches the floor when it tells one person which of two good things to choose on an ordinary morning.



Three Priorities on a Card

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER ONE

Twenty to nine, the card

It is twenty to nine on a Monday, and the claims office on the second floor is still half empty. On every desk there is a card the size of a postcard, left overnight by someone from internal communications, and the team leader of the household claims team is holding hers beside a cup of tea she has not yet tasted. The card is well made. It has the company's new name for its strategy across the top, and under it three priorities in a confident serif: *Members first. Simpler every day. Grow with care.* On the back, a line from the chief executive thanks everyone for the part they will play.

At five past nine the first claim of the week comes up on her screen. A burst pipe under a kitchen floor, a family with two small children, a request to settle today so the plumber can start tomorrow. The procedure says a claim of this kind goes to a loss adjuster first, which takes about a week. *Members first* says settle it. *Grow with care* says check it. *Simpler every day* could be read either way. She turns the card over, as if the answer might be on the back, and it is not. Her shoulders already know that whatever she decides will be the wrong one for somebody upstairs.

She is a composite, drawn from the way these mornings are described in the research below and in the house's own practice, and no firm is meant. Most people who have worked inside an organisation that has recently launched a strategy have held that card.

Ammanuel Santa Anna's canon has a sentence about plans of this kind that is uncomfortably exact about where they were written.[1] Writing of what governance produces when it does not see itself, it says that it "generates the kind of strategic plans that look brilliant on paper and crumble on contact with reality — because the reality they were designed for was the projected reality of the most powerful people in the room, not the actual reality the organization inhabits." The card was designed for a room on the sixth floor. The burst pipe is on the second.

This chapter makes one claim. A strategy has not reached the floor when it has been communicated; it reaches the floor when it tells one person which of two good things to choose on an ordinary morning. Until then it has been announced, and the work of deciding what it means has been handed, without anyone saying so, to the people least able to see the whole of it.

Communicated, not understood

In March 2015 three researchers published in the *Harvard Business Review* the findings of a long survey of managers inside large organisations.[2] Donald Sull, Rebecca Homkes and Charles Sull had asked, among many other things, how often senior leaders talked about strategy and whether the managers below them could say what it was. Most managers, as the study is reported, believed their leaders communicated the strategy often enough. Yet only about half of the middle managers asked could name even one of their company's top five priorities. The authors framed this as one of five myths about execution: that communicating is the same as understanding.

The evidence should be graded before it is leaned on. This is survey work published in a practitioner magazine, not a peer-reviewed study, and the figures here are taken from the article as it has been summarised, because the original text was not reachable for this book. Self-report about one's own knowledge of a strategy is a limited instrument. What the finding shows with some confidence is the shape of the gap: the people sending the strategy and the people receiving it describe the same communication very differently, and the senders are more satisfied with it.

It is not a new observation. In 1989 Donald Hambrick and Albert Cannella argued that implementation is a matter of substance and of selling, and that the selling has to go down, across and outward as well as up, to everyone whose daily choices the strategy depends on.[3] Their point was not that strategies need more slogans. It was that a strategy which is not understood well enough to argue with has not been sold at all; it has been distributed.

The card on the claims team leader's desk is distribution. It tells her the strategy exists, and that it matters to someone. It does not tell her what it asks of her at five past nine. Santa Anna's canon, describing who must be in the room when a value stream is mapped, puts the difficulty in one sentence that would serve as the first line of any deployment plan: "The executive who sees strategy, the middle manager who sees implementation, the frontline worker who sees daily reality — each holds a piece of the picture that the others cannot access from their vantage point." The card was written by the first, carried by the second and left on the desk of the third. Nobody asked the third what she would need to know.

There is a temptation, faced with findings like these, to blame the people who could not name the priorities. That reading has the direction wrong. When half of a capable management population cannot repeat what it has been told many times, the likeliest explanation is not that half of them were not listening. It is that what they were told did not attach to anything they do. A priority that never changes a decision is easy to forget, because nothing in the working day ever reminds you of it.

A map with arrows

The most influential attempt to close this gap came from Robert Kaplan and David Norton, who spent the 1990s arguing that a firm's measures should tell the story of its strategy. In a 1996 paper in *California Management Review* they set out how the balanced scorecard could be linked to strategy, with its financial, customer, internal-process and learning measures joined by chains of cause and effect, so that a person could in principle read from her own measure up to the strategy it served.[4] Their later book, *Strategy Maps* (2004), drew those chains as a single page of boxes and arrows.[5]

The house's Luminous 100, its reading of the frameworks every consultancy sells, is generous to the idea and notes the thing practitioners tend to find for themselves: "The strategy map connecting the four perspectives is often more valuable than the scorecard itself." The map is valuable because it forces a leadership team to say, in public, what it believes causes what. It replaces *members first* with something closer to an argument: if claims are settled faster, members renew; if members renew, the cost of winning new ones falls.

The same reading carries the map one step further, in its own terms. Living-systems awareness, it says, "transforms the strategy map from a linear cause-and-effect chain into a living circulatory system". Read plainly, that is a warning about where the arrows stop. On most maps they stop at the department, because the department is the smallest unit the people drawing the map can see. There is a box for claims and an

arrow from it to member retention. There is no arrow to the team leader's screen at five past nine, and that is the arrow that decides whether the rest of the map is true.

This is not a criticism of Kaplan and Norton's tool. It is a description of how tools are used. A map drawn at the top and handed down is still a message from the sixth floor; the arrows are the leadership team's beliefs, and they are only tested when somebody on the second floor makes a choice and the result comes back. If nobody on the second floor knows which arrow her choice belongs to, the map cannot learn. It becomes a picture of a strategy rather than an instrument for running one.



What the card leaves out

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER ONE

What the card leaves out

Look again at the three priorities. Each of them is something nobody could disagree with, and that is the trouble. A strategy is a set of choices, and a choice has a cost: it says what the firm will do when two good things conflict. *Members first* and *grow with care* conflict every time a claim could be settled quickly or checked carefully, which in a claims office is most of the day. The card names both goods and leaves the conflict to the reader.

The historian and consultant Stephen Bungay, in *The Art of Action* (2011), traced this problem to the Prussian army of the nineteenth century, and described three gaps that open between a plan and its results: between what leaders would like to know and what they know, between what they want people to do and what people actually do, and between what they expect actions to achieve and what those actions achieve.[6] His argument, summarised here from the book and not reopened for this book, is that the usual response to each gap, more detail, more instruction, more control, widens the next one. The army's answer was to pass down intent rather than instructions: what is to be achieved and why, and the freedom to choose how. A subordinate who knows the intent can decide well when the situation changes, which it always does.

Santa Anna's canon arrives at the same principle from the other direction, asking where the intelligence for a decision actually lives: "If the intelligence lives on the frontline — as it does for most operational decisions, customer interactions, and day-to-day process adjustments — then the frontline should make the decision." The burst pipe is exactly such a decision. The team leader knows the family, the plumber, the photographs on the file and the eleven similar claims she settled last month. What she lacks is not judgement. It is the intent: which of the three priorities governs when they conflict, and how much risk the firm is prepared to carry to keep a promise to a member.

The canon has a name for a priority that floats free of that kind of decision. Writing of the room every organisation keeps without a door, it lists among the exhibits "The fact that the values on the wall bear no resemblance to the values in the hallway." A card on a desk is a small wall. If the team leader settles the claim and is later asked why she skipped the adjuster, she will learn what the value in the hallway is, and so will the eleven people who sit near her.

There is a cost here, and it is worth naming plainly because it is usually invisible. Every team leader in the building is resolving the same conflict on the same Monday, separately, without telling each other. Some will settle; some will send everything to the adjuster to be safe. The firm will pay for both: in members who leave because they waited a week, and in claims that should have been checked. Each of those costs is foreseeable, and none of them appears in any report as the price of an unclear strategy. They appear as claims performance, as retention, as leakage, each owned by a different department. The difference between a strategy decided once, clearly, and the same strategy decided nine ways on the second floor is real money, and nobody measures it. How it could be measured, [honestly enough that both sides would sign the number]({{BOOK:orgdev-flow}}), is a question the house has answered elsewhere.

Why the card says so little

It is worth asking why the card is written the way it is, because the people who wrote it are not careless. The three priorities were agreed in a room of eight executives over two away-days, and each word was argued over. *Members first* was the chief operating officer's phrase. *Grow with care* was the finance director's, added late, after a long afternoon about the cost of claims. The two of them disagreed, at the table, about whether the firm should settle more claims quickly or check more of them carefully, and the disagreement was resolved in the way such disagreements often are: by keeping both phrases.

That is the hidden mechanism behind most strategies that arrive as three good words. The vagueness is not a failure of drafting. It is the residue of a choice that the leadership team did not finish making, written down in a form that allows each executive to read it as a victory. The card is the minutes of a truce. It travels down the building intact, and then it lands on the second floor, where the truce cannot hold, because the team leader has to settle the claim or send it to the adjuster, and there is no third option.

Santa Anna's canon describes what a room looks like when it agrees in this way. Writing about the quiet signs that a group has stopped engaging, it lists among them "Compliance without commitment. Agreements that carry no energy. Head nods that are automatic rather than genuine." The away-day ended in agreement. The agreement carried no decision, and so it carried no energy down the building either.

The cost of the truce is not paid by the people who made it. The chief operating officer and the finance director will each, in the coming quarter, see evidence that their reading was right: faster settlements in one report, lower leakage in another. Both can be true, because nine team leaders are resolving the conflict in nine directions, and each executive will see the teams that agreed with them. The leadership team will not see the conflict at all, because the place it lands is two floors down, at five past nine, in the stomach of the person holding the card.

There is a practical consequence for anyone who writes a strategy. The test of a priority is not whether everyone in the leadership team can agree with it. It is whether the leadership team, asked what the priority means for a burst pipe on a Monday, gives the same answer. If eight executives give two answers, the card is not ready to leave the room.

If, then, on Monday

Psychology has a finding that bears directly on the card, and it is one of the better replicated results in the study of goals. In 1999 Peter Gollwitzer described what he called implementation intentions: plans that specify in advance when, where and how a goal will be acted on, in the form *if situation Y arises, then I will do X*.[7] A goal held as an intention, *I will put members first*, is often not acted on when the moment comes, because the moment does not announce itself. An if-then plan links the goal to the situation, so that the situation itself becomes the cue. In 2006 Gollwitzer and Paschal Sheeran gathered ninety-four independent

tests and found that forming implementation intentions had a positive effect on reaching goals of medium-to-large size.[8]

The research is about individuals and their own goals, mostly in studies of health, study habits and similar tasks, and it should not be stretched into a claim about organisations. But the mechanism translates cleanly. *Members first* is a goal intention. *If a member asks for a small claim to be settled the same day and the photographs show the damage clearly, then I settle it and note why* is an implementation intention. The second can be acted on at five past nine. The first can only be agreed with.

The house has written about how [an obstacle written down in advance arrives as an appointment rather than a verdict]({{BOOK:rsd-rsd-and-self-confidence}}), and the same holds at work. A strategy that has been turned into three or four if-then sentences for each team has had its conflicts settled before the morning they arrive. A strategy left as three good words has had its conflicts postponed to the moment they are hardest to settle well, by the person least able to see their consequences.

Santa Anna's canon describes what a phrase looks like when it does reach the floor, in a composite from its chapters on organisational soul. A research company had stopped saying the sentence its founders used, *for the patients we haven't met yet*, and when it came back, it came back not as a campaign. The canon is precise about the difference: it became "not a slogan, not a brand, but a living commitment that people used to orient their decisions". The test of a strategy on the floor is the same. It is whether anyone uses it to decide something.



The back-brief

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER ONE

The back-brief

What changes this is modest, and it begins one level up. Picture the area manager who sits above the claims team leader and eight others like her. In the week after the cards arrive, she does one thing differently. Instead of a briefing, she asks for a back-brief, a practice Bungay describes from the military and which many firms already use under other names: each team leader writes, on a single page and in her own words, what the strategy asks of her team, and sends it back up.

The page has four short parts. What the strategy asks of this team, in a sentence a new starter would understand. The one or two decisions it changes, written as if-then sentences, using a real case from last month. What the team will stop doing, or do less of, because of it. And what the team leader needs from above to do it: an authority, a piece of information, a decision only someone else can make. The claims team leader writes that members first means the same-day settlement of clear, small claims, and that she needs to be told how large a claim she may settle without an adjuster, and that she cannot answer that herself.

The order of what happens next matters more than the page. The area manager reads all nine back-briefs before she meets anyone. She finds that four team leaders have read *members first* as settle-quickly, three as be-kind-on-the-phone, and two have not changed anything. That is not a failure of the team leaders. It is the most useful information the strategy has produced since it was launched, because it shows exactly where the card left a choice open. She takes the nine pages, unedited, to her own director, and asks the question none of them can answer: which priority governs when they conflict. Within the week the answer comes back down as two sentences, and each team leader rewrites her if-then lines against it.

The canon's chapter on the keystone effect describes the person the area manager has become, in a question it asks of every organisation: "Who translates executive decisions into meaning for the front line?" The answer in most firms is someone doing it informally, at the coffee machine, without authority. The back-brief makes the translation formal, visible and answerable. It turns the middle of the organisation from a relay into a place where the strategy is tested before it is spent.

There are failure modes, and they are predictable. A back-brief that is graded, praised or corrected for its wording teaches team leaders to write what the area manager wants to read. A back-brief that asks for something and receives no answer teaches them not to ask. And a back-brief demanded of team leaders by an area manager who has not written one herself, upward, about the strategy she received, is asking for a candour she has not shown. Chapter 7 carries the run sheet for this practice, and Chapter 6 places it in the order of a first ninety days. What can be said here is that the first round of back-briefs usually disagrees with itself, and that the disagreement is the point.

The next place the plan is lost is the person the area manager has just become, the manager in the middle who is handed a target and has to pass it on (Chapter 2), and further down the same road lies the number that starts to stand in for the purpose it was meant to serve (Chapter 3).

Writing the back-brief in your own building

The article this chapter began as ended at the team leader's desk. The book can go further, into the finer points a leader meets when she tries the back-brief in her own organisation, and the first of them is knowing whether her strategy has reached the floor at all.

There is a simple way to find out, and it needs no survey. Choose ten people, at least half of them within two levels of the work itself, and ask each one, privately and with no notice, one question: *what does the strategy ask you to do differently on Monday?* Write down the answers word for word. Most leaders who do this find three kinds of answer. Some people repeat the words on the card, accurately and without content. Some describe something their own manager has told them, which may or may not be the strategy. And a few describe a real choice they now make differently, in their own words, with an example. The proportion of the third kind is the nearest thing a leader has to a measure of whether her strategy has been deployed or only distributed. It is not a statistic and should not be reported as one; it is a reading, taken by hand, and it is usually sobering.

The second finer point is that the back-brief begins at the top, not in the middle. Before any team leader is asked to write one, the leadership team should be asked the question the claims team leader could not answer: for each priority, take one ordinary case from last month, and say what the priority means for it. If the leadership team gives two answers, the strategy is still a truce, and every back-brief written below it will expose the truce rather than resolve it. That is useful, but it is cheaper to find out in the boardroom. The practical rule is that no priority goes on a card until the people who wrote it have answered the same case the same way.

The third is what the back-brief page should and should not contain. It should contain the team's own words, a real case, an if-then sentence or two, what will stop, and what is needed from above. It should not contain a score, a plan with milestones or a request for budget, because each of those turns the page into a negotiation, and a negotiation is written for its reader rather than for its truth. A page of about two hundred words is enough. A page of two thousand has been written for the area manager.

The fourth is the reading. The person who receives the back-briefs reads all of them before replying to any, because the value is in the comparison. Four team leaders reading *members first* one way and three another is not nine opinions; it is one finding about the card. The reader sorts the pages by the choice each makes, notes where they agree and disagree, and takes the disagreements upward unedited, with the case attached. She answers every page within the week, even if the answer is only that the question has gone up and when it will come back.

The fifth is the stopping. Strategies are almost always written as additions, and the back-brief's line on what will stop is the one teams find hardest to fill in. It is also the line that most changes Monday. A team that has written *we will stop sending every water claim to the adjuster* has made the strategy real in a way that no amount of new activity can. If the line is empty on most pages, the strategy has asked for more without permitting less, and the people above should say what can be dropped before they ask for anything else.

The sixth is the leader's own back-brief. The area manager who asks nine team leaders to write one should write her own first, upward, to her director, about the strategy as she received it. It is the same principle Chapter 2 develops at length: the level above has to be willing to be caught out by a question it cannot an-

swer. A back-brief practice in which only the bottom writes is an inspection. One in which every level writes upward is a conversation about the strategy, held in writing, at the speed of a week.

And the decision that most changes the result is what happens to the first back-brief that disagrees with the strategy. If the team leader who writes that *members first* cannot be done without more authority is thanked, and given an answer, and sees the card's meaning change because of what she wrote, the second round of back-briefs will be candid. If she is told she has misunderstood the strategy, the second round will be accurate copies of the card. Chapter 6 places the back-brief in the first month of a ninety-day plan, after the leadership team has tested its own priorities, and Chapter 7 gives the floor reading, the truce test and the back-brief page as working pages.

Twenty to nine, a month later

It is twenty to nine on a Monday in the following month, and the card is still on the desk, a little curled at one corner. Under the three printed priorities, in pencil, the team leader has written two lines of her own. The first says which claims she settles on the day. The second says who she calls when she is not sure.

At five past nine the first claim comes up. A leaking roof, a photograph of a stained ceiling, a request to settle this week. She reads it once, checks it against the second line, and settles it before the tea has gone cold.

The card has not changed. It is still three good words on a postcard, and now it is a strategy.

CHAPTER TWO

The Ball Thrown Back

The middle is where a target is most often lost, because the only direction it is allowed to travel is down.



The Ball Thrown Back

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER TWO

Ten to seven, the car park

It is ten to seven on a Thursday evening in November, and the regional manager is sitting in her car in the car park of the third hotel she has visited today. The engine is off. The windscreen has begun to mist at the edges, and the light from the hotel's lobby falls in a long yellow square across the tarmac beside her. On her phone is the email she has been waiting for since September: next year's targets for her twelve hotels, one line each, agreed by the executive team that afternoon. She is asked to cascade them to her general managers by Monday.

She reads the list twice. Ten of the lines are hard and fair. Two are not. Both belong to hotels where she spent an hour this week, and at both the general manager walked her past a housekeeping rota that already runs two people short on a Saturday. The only way either hotel reaches its line is by cutting the hours that keep its rooms clean, and the strategy deck that came with the email says, on its second page, that next year is the year of the guest. Her stomach has the particular heaviness of a person who knows she will be the one to say the number out loud on Monday, to people she respects, and who does not believe it.

She is a composite, drawn from the way these evenings are described in the research below and in the house's own practice, and no firm is meant. Anyone who has managed managers has sat in that car.

Ammanuel Santa Anna's canon describes her position with more precision than most management textbooks manage.[9] Writing about whose view of an organisation can be trusted, it says:

> *Middle managers occupy the most paradoxical position: they see enough to know that the view from the top does not match the reality at the bottom, but they often lack the authority to act on this knowledge and may be incentivized to smooth the discrepancy rather than surface it.*

This chapter makes one claim. The middle of an organisation is where a strategy is most often lost, and it is lost not because middle managers resist it but because the only direction most firms let a target travel is down. A strategy reaches the floor intact when the person in the middle is allowed, and expected, to throw it back with the truth attached.

Counter-effort

There is a familiar story about middle managers and strategy, and the research supports part of it. In 1986 William Guth and Ian MacMillan studied how middle managers responded to strategies they believed compromised their own interests.[10] Working from expectancy theory, they found strong if indirect evidence that such managers could redirect a strategy, delay it, reduce the quality of its implementation, or sabotage it outright. They gave the behaviour a useful name: counter-effort. It is the energy a person spends, quietly, making a decision she disagrees with turn out the way she expected.

The evidence is survey-based and the authors called it indirect, and it should be read that way. It has lasted because most executives recognise the pattern, and because it confirms what they already suspect: that the middle is where good strategies go to be slowed down.

What the study is less often read for is its cause. Counter-effort, in Guth and MacMillan's account, follows from a belief that one's interests are being compromised and that there is no other way to say so. The regional manager in the car park is not plotting. She is facing a choice between two kinds of counter-effort. She can pass the two lines down as they stand and watch the general managers find the hours somewhere the strategy will not see, or she can quietly soften them, adjusting what she asks of the two hotels and car-

rying the difference herself, and hope nobody notices in March. Either way the executive team will not hear what she knows. Santa Anna's canon describes the second choice with some tenderness: "The middle manager who carries the emotional burden of translating corporate euphemisms into human reality for their team, cushioning the impact of decisions they had no voice in making."

The canon also describes where the cushioning ends up. Writing about information in hierarchies, it observes that each layer filters what it passes upward to match what the next layer expects, and then states the result in a single line: "By the time a frontline reality reaches the C-suite, it may be unrecognizable." The two housekeeping rotas will reach the executive team in the spring as a variance on a line, explained in a sentence about labour markets.

Up, as well as down

Through the 1990s two researchers, Steven Floyd and Bill Wooldridge, set out to describe what middle managers contribute to strategy beyond carrying it out. Their answer was that the contribution runs in both directions. In a 1992 research note they distinguished between the influence middle managers exert downward, implementing and facilitating, and the influence they exert upward, championing alternatives and synthesising information for those above them.[11] In 1997, studying 259 middle managers, they found that managers in boundary-spanning roles reported more strategic influence than others, and that firm performance was associated with more uniform downward influence and more varied upward influence across the middle-management cohort.[12]

The second finding deserves a careful reading. It is correlational, from self-report, and it does not show that upward influence causes better performance. What it does show is that in the better-performing firms in their sample, middle managers did not all push the same message upward. Some of them, at least, were sending back something different from what had been sent down. A firm whose middle managers all report that the targets are achievable is not, on this evidence, a firm with an aligned middle. It may be a firm with a quiet one.

A review of the field by Wooldridge, Schmid and Floyd in 2008 drew the same picture across two decades of work: middle managers as interpreters, sellers and synthesisers of strategy, not merely its conduits.[13] And Frances Westley, writing in 1990 about middle managers in bureaucratic organisations, described the particular moment at which this contribution is won or lost.[14] It is the strategic conversation between a manager and the person above her. Middle managers who are excluded from such conversations, she argued, become de-energised, and the exclusion is inefficient and in the end expensive. But inclusion alone does not guarantee energy: managers can be invited into the conversation and still leave it flat, if what they bring is not taken up.

Read the car park with those findings in mind. The regional manager has not been excluded from strategy; she has been sent it. She has, as Floyd and Wooldridge would put it, a boundary-spanning role: she walks

the hotels and attends the executive meetings. She holds precisely the information upward influence exists to carry. What she lacks is a channel through which a target can travel back up carrying the housekeeping rota. Without one, her knowledge becomes counter-effort or cushioning, and either way it is lost.



The work of holding two things

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER TWO

The work of holding two things

There is a cost to the middle that does not appear in the research on strategy at all, and it is the cost the regional manager feels first. In a three-year field study of a large firm attempting radical change, Quy Huy

found that the middle managers who helped their groups adapt did two seemingly opposed things at once.[15] They committed emotionally to change projects they personally championed, and they attended carefully to the emotions of the people the change landed on. Low commitment to change produced inertia. High commitment with little attention to the people produced chaos. Doing both, which Huy called emotional balancing, was what allowed the organisation to change and still keep serving its customers well.

Huy's study is one firm, studied closely, and it should not be read as a law. But it describes something many managers will recognise as their working week. To pass the two lines down and also look the general managers in the eye on Monday, the regional manager must believe in next year's direction and also honour what the housekeeping rota is telling her. Holding both is work. It is done by the people a firm most depends on, and nobody budgets for it.

Santa Anna's canon, writing about what happens when promises are made under pressure, names the cost of not doing that work honestly. A commitment extracted under pressure, it says, "is not a genuine commitment. It is a performance of commitment, and the person making it knows the difference even if no one acknowledges it." The general managers who accept their lines on Monday without argument will have performed commitment. In March the performance will end.

The canon also describes, in the same chapter, the shape of the culture this produces when it is repeated year after year: "The result is a culture of broken promises that everyone expects and no one names." Every January the targets go down. Every December some of them are missed, by the margins the middle predicted. The board learns that the organisation is weak at execution. The organisation learns that the targets are not really meant.

Sense made sideways

If the middle cannot send what it knows upward, it does not stop knowing it. It talks sideways. In 2004 Julia Balogun and Gerry Johnson published a study of middle managers living through the restructuring of a large organisation, followed over many months through diaries and interviews.[16] Their finding, summarised here from the published study and not reopened for this book, was that the managers made sense of the change chiefly with one another, in informal conversations among peers, and largely in the absence of the senior managers who had designed it. What the restructuring came to mean on the ground was settled in those conversations, and it was not always what had been intended.

That is the regional manager's Friday, if nothing changes. She will ring the two other regional managers she trusts on the drive home. Each of them will have their own two lines that will not hold. Between them, in three phone calls, they will agree what next year's targets really mean, which ones are serious, which ones will be quietly adjusted, and what to say to the general managers so that nobody panics. It is the same corridor the house has described in [the meeting after the meeting]({{BOOK:orgdev-flow}}), moved into a car and spread across a region. It is intelligent, and it is invisible to the people who set the targets.

The sideways conversation is not a failure of loyalty. It is the system doing what living systems do when a channel is blocked: finding another one. The canon names the general pattern in its chapter on information, in the line already quoted, and the practical consequence follows from it. A firm that wants to know what its middle knows has to give the knowledge a road upward that is quicker and safer than the phone call on the drive home. Otherwise the phone call wins, every November.

A ball, and a hand to catch it

There is a long-standing management practice built for exactly this, and it comes from Japan. Hoshin kanri, often translated as policy deployment, grew out of Japanese quality management in the 1960s and was set out in English by Yoji Akao in a book first published in 1991.[17] At its centre is a process the Japanese practitioners called catchball. A goal is not handed down. It is thrown, with its reasons, to the level below, which considers what it would take, what it can offer and what it would need, and throws it back. The ball goes back and forth until both levels hold a target they believe, and the plan to reach it. Alongside it runs nemawashi, a gardener's word for preparing the roots before a tree is moved: the patient, one-to-one conversations in which a proposal is tested with the people it will affect before anyone is asked to decide.

Barry Witcher and Rosie Butterworth studied hoshin kanri in practice over two years in three Japanese-owned manufacturing subsidiaries in Britain, and published their findings in 2001.[18] They described it as a framework that links the achievement of top management goals with daily management at the operating level, translating policy objectives into targets for quality, cost, delivery and education. They found it participative in how objectives are developed and deployed, and driven by a process of review; they also found it hard to sustain through changes in people and organisation, and difficult to administer across departments. In the same year Ceri Tennant and Paul Roberts published a study of how the catchball process itself is implemented.[19] The research is case work in a small number of firms, and much of the wider writing on hoshin kanri comes from practitioners and consultants. What it supports is the description of the practice and its difficulties, not a measured effect on results.

What catchball changes is the direction a target is allowed to travel. The house's Luminous 100 makes the same point about the most popular modern goal framework, objectives and key results, in its own language: it says that living-systems awareness "reveals that cascading OKRs down the org chart is a mechanistic practice that fragments the living whole". A target that only falls cannot learn anything on the way down. A target that is thrown and returned carries information in both directions, and the return is where the information is.



The car park, on Monday

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER TWO

The car park, on Monday

Picture the regional manager doing it differently. On Monday morning she does not cascade the twelve lines. She sends each general manager the line for their hotel with the two pages of the strategy deck that explain it, and asks three questions, to be answered by Wednesday on a single page: what it would take to reach this line with what you have now; what you could reach if one thing changed, and what that thing is; and what reaching it would cost elsewhere in the hotel, in your own words.

Ten of the pages come back saying, in effect, hard but possible, with a list of what they will try. The two she worried about come back saying what she knew in the car park, and something she did not. One general manager has written that he could reach his line without cutting housekeeping if the group let him close the hotel's restaurant on Sunday evenings, when it serves eleven covers. The other has written, simply, that her line cannot be reached without reducing room cleaning, and that her guest scores will fall.

The order of what she does next matters more than the questions. She reads all twelve before speaking to anyone. She answers each general manager in writing within a day, even if the answer is only *received, and I am taking it up*. She takes the two pages, unedited, to her director, together with a page of her own: what the region can commit to, what it could reach if two decisions were made above her, and what the remaining gap would cost the guest. She does not soften the two lines herself. She throws the ball back, with the truth attached, and asks for a decision.

The canon's chapter on commitments sets out the discipline she is practising, in a phrase that could be printed above every catchball: "Commitments are calibrated to capacity." It is not a refusal of ambition. It is an insistence that a target agreed in November should still be believed in March, and that the place to discover it will not be is the car park, not the year-end review.

There are failure modes, and they are predictable. A director who answers the returned ball by restating the target has taught the region that the questions were decorative. A regional manager who sends the questions and then argues with the answers has done the same. And a firm that runs catchball only on targets it has already announced to the board has made it a negotiation over how to absorb the shortfall, not a way of setting the target. Chapter 7 carries the run sheet for this practice, and Chapter 6 places it in the order of a first ninety days.

The gain in this is easy to underrate because it shows up as an absence. When the ball comes back in November, the two hotels' shortfall is met as a decision about a restaurant and a rota, made by people who can see the whole region. When it does not, the same shortfall is met in March, as a variance, by people trying to explain it. The difference between the two is the cost of a strategy learning in the wrong month, and it can be measured from records the firm already keeps, which is how the house has argued [a gain should be measured before anyone shares in it]({{BOOK:orgdev-flow}}). The regional manager's page is the earliest point at which that difference is visible. Further down, the same line on her email will become a number that starts to stand in for the purpose it was meant to serve (Chapter 3), and the earlier problem, a strategy that arrives as three good words and no choice (Chapter 1), is the reason her director's answer has to come back as a sentence and not a slogan.

What the director owes

Catchball is usually described from below, as a way for the middle to speak. It is at least as much a discipline for the level above, because a ball that is thrown back has to be caught. The regional manager's director

now holds twelve pages and a summary, and a choice about what to do with them that is harder than it looks.

The easy response is to thank the region and keep the targets, on the grounds that the targets were agreed by the executive team and the board has seen them. That response is common, and it is the end of catchball in that firm, because every regional manager will hear about it within a week. The harder response is to treat the returned ball as information about the strategy, not as a negotiation over it: to ask what the two hotels' pages say about the year of the guest, whether the targets and the strategy can both be true, and, if not, which one gives way.

The canon's chapter on commitments sets out the questions a body should ask before it accepts any promise, and they read as if written for the director's desk: *"Do you genuinely have the capacity to fulfill this? What would need to be true for this commitment to be realistic? What support do you need? And what are you willing to let go of to make space for this?"* The last question is the one most often left out. A target that adds to the work and removes nothing is a target that will be met by removing something the target does not see, and in a hotel that is usually the housekeeping hours.

There is also a matter of time. Catchball takes longer than cascading. The targets will reach the general managers a week later than planned, and the executive team will have to meet once more before the year begins. That week is the cheapest week of the year in which to discover that two of the lines are wrong. Every later week is more expensive, and the most expensive is the one in March when the shortfall is discovered as a fact.

Finally, the director owes the region an answer in the same form the region used: in writing, specific, and quickly. A ball thrown back and caught in silence teaches the same lesson as a ball that is never caught. Witcher and Butterworth's finding that hoshin kanri is hard to sustain through changes in people is, read from the director's chair, a finding about this moment. The practice lasts as long as the people above keep catching.

Running catchball in your own building

The article this chapter began as followed one regional manager and one set of targets. The book can go further, into the finer points of running catchball across an organisation, and the first of them is choosing where to begin.

Catchball is best begun with the next cycle of targets rather than the current one. Targets already announced to a board carry a weight that makes every returned ball look like a request for relief, and the conversation becomes a negotiation over a shortfall. Targets not yet set can be thrown and returned honestly, because nobody has yet promised anything. In practice this means the first catchball is held in the months before the planning cycle closes, and the leadership team agrees in advance that the returned balls will be

read before the targets are fixed. If the cycle has already closed, choose one priority whose targets are still open, or one region, and begin there.

The second finer point is what travels with the ball. A target thrown down without its reasons cannot be returned with anything useful, because the level below can only say whether the number is achievable, not whether it is the right number. Each target should go down with the part of the strategy it serves, in a sentence, and with what the level above believes about how it could be reached. The level below then answers three questions: what it would take to reach the target with what is there now; what could be reached if one thing changed, and what that thing is; and what reaching it would cost elsewhere. The third question is the one that finds the housekeeping rota. It is also the one people are most reluctant to answer in writing, and the reply to the first honest answer to it decides whether it is ever answered again.

The third is the shape of the return. A returned ball is most useful when it separates three things: what the level below commits to, what it could reach with a named decision from above, and what it believes cannot be reached without a cost to something else the strategy values. Mixing the three produces a single, softened number, which is the cushioning described earlier in this chapter. Keeping them apart lets the level above see exactly which decisions are its own.

The fourth is the number of rounds. In the hoshin kanri literature the ball goes back and forth until both levels hold a target they believe. In most organisations two rounds are enough for most targets and a third for one or two. More than three usually means that the level above has not yet decided what it is willing to let go of, and the right move is to stop throwing and make that decision.

The fifth is the people in the middle of the middle. Catchball between a director and a regional manager does little if the regional manager then cascades to her general managers in the old way. The practice has to run at every level the target crosses, and each level has to hold its own returned balls before passing anything upward. That takes longer than a cascade, by a week or two at each level. The canon's line on commitments calibrated to capacity is the reason the time is worth spending: a target that has been tested at every level is one that every level has reason to believe.

The sixth is the sideways conversation. The regional managers will still talk to each other on the drive home, and they should. The practical step is to bring that conversation inside: before returning their balls, the regional managers meet for an hour, without their director, to compare what their general managers have said. They will find the same patterns, and they will return balls that say so. A pattern reported by six regions is harder to dismiss than one reported by one, and it is also more often right.

And the decision that most changes the result is what the director does with the first returned ball that says a target is wrong. If it is answered with a decision, in writing, within the week, and the target or the strategy changes because of it, the practice will hold. If it is answered with the original target restated, catchball is over, and every later round will be a ritual. Chapter 6 places catchball in the first two months

of a ninety-day plan, for one priority, and Chapter 7 gives the three questions, the return page and the director's questions as working pages.

Ten to seven, a week later

It is ten to seven on the following Thursday, and the regional manager is in the same car park. The windscreen is misting again at the edges. On her phone is an email from her director, sent at four. It says that the Sunday restaurant can close for a trial of three months, that the second hotel's line has been lowered and the difference moved to the group's own budget for the year, and that the executive team would like the three questions used across every region next November.

She reads it once. Then she writes two short replies, one to each general manager, and starts the engine.

The light from the lobby still falls across the tarmac in the same yellow square. She drives out through it.

CHAPTER THREE

The Number That Stood In

A number set alone, without its purpose and its pair, slowly takes the strategy's place.



The Number That Stood In

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER THREE

A quarter past four, the screen

It is a quarter past four on a Friday, and the screen on the wall of the contact centre has just turned green. It shows one number for each team, the average time an adviser spends on a call, and for the first time this quarter the figure for her team is under the line. The team leader is standing at the end of the row with a lanyard in one hand and a list of the week's escalations in the other. Somebody further down the room claps, once, softly, and then stops, as if unsure whether it is allowed.

Two desks away her most experienced adviser is still on a call. It has been running for thirty-eight minutes. On the other end is a man in his eighties whose broadband has failed for the third time in a month, and the adviser is talking him, slowly and kindly, through the back of his router, one cable at a time. At the start of the quarter the company's objective was printed at the top of every team's plan: *customers who never have to call twice*. The key result under it was the handling-time line. The team leader looks at the screen, then at the adviser, then at the escalation list, where she knows that four of this week's calls were transferred to another team at the nine-minute mark, and her chest goes tight in the particular way it does when two things she is responsible for point in opposite directions.

She is a composite, drawn from the way these afternoons are described in the research below and in the house's own practice, and no firm is meant. Most people who have managed to a number have watched it turn green for the wrong reason.

Ammanuel Santa Anna's canon has a sentence for the moment, written about a firm that measured everything well and was losing what mattered most.[20] "When you can only count certain things," it says, "those things become the things that count." The adviser's thirty-eight minutes will count against the team. The man in his eighties, who will not call again this month, will not be counted at all.

This chapter makes one claim. A strategy reaches the floor as a number, and that is right; but a number set alone, without the purpose beside it and without a second number to watch it, slowly takes the strategy's place. The gain lost here is not the gap between the target and the result. It is the gap between the result and the purpose, and most firms have no instrument that sees it.

Goals work

It would be easy, from the contact centre, to conclude that targets are the problem. The evidence says otherwise, and it deserves to be stated fairly first. In 2002 Edwin Locke and Gary Latham summarised thirty-five years of research on goal setting in the *American Psychologist*.[21] The core of it is one of the most replicated findings in the study of work: specific, difficult goals lead to higher performance than vague goals or the instruction to do one's best, across hundreds of studies, many settings and many countries. Goals direct attention, raise effort, increase persistence and prompt people to look for better ways of working. The authors were also explicit about the conditions: people need to be committed to the goal, to receive feedback on progress, to have the ability to reach it, and to face a task that does not overwhelm them.

That last condition matters for the contact centre, and Locke and Latham wrote about it themselves. When a task is complex and people do not yet know how to do it well, a demanding outcome goal can make performance worse, because it pushes people to try harder at a method that does not work. In 2005 Gerard Seijts and Latham set out the remedy: in such situations, set a specific, challenging learning goal instead, a goal to discover a number of effective ways of doing the work, and set the outcome goal later.[22] A learning goal draws attention away from the end result and toward the discovery of how to reach it.

So the theory, read whole, does not say *set a number and push*. It says set a number when people know how to reach it, and something else when they do not, and watch for what the number does not see. The contact centre's key result was set as if helping the man in his eighties were a simple task with a known method. It is not.

A warning label

In 2009 four researchers, Lisa Ordóñez, Maurice Schweitzer, Adam Galinsky and Max Bazerman, published an essay with a title designed to be remembered: *Goals Gone Wild*.[23] They argued that the benefits of goal setting had been overstated and its harms largely ignored, and they listed the side effects they saw: a narrow focus that neglects everything outside the goal, distorted attitudes to risk, a rise in unethical behaviour, inhibited learning, corrosion of culture and reduced intrinsic motivation. Goal setting, they wrote, should be treated less like an over-the-counter remedy and more like a prescription-strength medication, with careful dosing, close supervision and a warning label.

Locke and Latham replied in the same issue of the journal, and sharply.[24] They accused the critics of drawing their central theme from anecdote rather than research, of citing the literature unrepresentatively, of misreporting results and of using emotionally laden metaphors, and they pointed out that almost every pitfall listed had already been described in their own work. Both papers are essays in a management journal, not new studies, and a reader should hold the dispute as what it is: a disagreement between serious scholars about emphasis, not about the basic facts.

One of the studies at the centre of it is worth describing, because it is a careful experiment and its finding is specific. In 2004 Schweitzer, Ordóñez and Bambi Douma gave people a task in a laboratory and either a specific goal or the instruction to do their best, and then gave them the chance to misreport their own performance.[25] People with unmet goals were more likely to overstate what they had done than people asked to do their best, with or without a financial reward attached, and the effect was strongest among those who had fallen just short. It is one laboratory study, with students, on a simple task. What it shows is not that goals make people dishonest. It shows where the pressure gathers: at the line, among the people closest to it.

The four transfers at the nine-minute mark sit exactly there. Nobody in the team leader's row is lying. They are doing what the line rewards, at the point where the line is closest, and the canon describes the pattern without any accusation at all. Writing about why fear-based accountability produces the behaviour it is meant to prevent, it says that people "invest energy in managing the *appearance* of performance rather than the reality of it — producing impressive reports, hitting narrow metrics, and performing competence while the actual work suffers from the attention diverted to self-protection."



Targets and terror

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER THREE

Targets and terror

The largest natural experiment in governing by numbers took place in the English public health service in the 2000s, and it was studied closely. In 2006 Gwyn Bevan and Christopher Hood published an analysis with a title that is also a summary: *What's measured is what matters*.[26] The government of the day had combined targets with what they called an element of terror, published rankings and real consequences for managers who missed. Bevan and Hood drew the parallel with the Soviet planning regime, which was initially successful and then collapsed, and they named the two assumptions such a system rests on. The first

is synecdoche, taking a part to stand for the whole: a waiting time standing for the quality of care. The second is that problems of measurement and gaming do not matter. Their evidence, drawn from reported successes, measurement problems and documented gaming, showed that both assumptions were fragile.

The study is careful to say that targets produced real improvements as well as gaming, and it should not be read as an argument against measurement. Its lasting contribution is the word synecdoche, because it names exactly what happens on the contact centre's wall. The handling-time figure is a part. The objective, customers who never have to call twice, is the whole. For the first weeks of the quarter the part and the whole moved together, because shorter calls often meant better-trained advisers. Then they separated, and the part kept going.

Steven Kerr had described the general case in 1975, in a paper whose title has entered the language: *On the folly of rewarding A, while hoping for B*.[27] Organisations, he showed with examples from many settings, routinely reward one behaviour while hoping for another, and are then surprised to get the behaviour they rewarded. Santa Anna's canon, describing a technology firm that had invested heavily in a sophisticated goal framework, puts the end state in three short sentences: "Meridian had optimized itself into hollowness. The metrics were excellent. The soul was departing."

What the number cannot carry

There is a reason the separation is so hard to see from above, and the canon is precise about it. Writing of the framework the contact centre was using, objectives and key results, it gives the method its due and then names its limit: it "reinforced the assumption that what matters can be expressed as a measurable outcome within a defined timeframe — an assumption that works beautifully for product launches and revenue targets but poorly for relational trust, ethical development, creative incubation, and the slow, non-linear work of organizational culture change." A thirty-eight-minute call with a man in his eighties is relational trust. The method can count its length. It cannot count what it bought.

The practitioners who popularised the framework knew this, and said so. Andrew Grove, who built the method at Intel and described it in *High Output Management* (1983), argued that indicators should be paired, so that a measure of an effect is always watched beside a measure of its counter-effect: a measure of speed beside a measure of quality, a measure of stock levels beside a measure of shortages.[28] John Doerr, who carried the method to Google and wrote *Measure What Matters* (2018), drew a distinction between committed key results, which a team is expected to meet in full, and aspirational ones, set high on purpose, where falling short is expected and is not a failure.[29] Both points are summarised here from the books, not reopened for this book. Both are routinely lost when the method is adopted, because the single number fits on a screen and the pair does not.

The house has written about why [a number set beside a name is heard as a verdict]({{BOOK:rsd-master-rsd-in-acedemia-final-manuscriptt}}), and about the trouble with [a goal asked to carry a person's worth]

({{BOOK:rsd-rsd-and-self-confidence}}). The contact centre adds the organisational version. A key result asked to carry the whole of a strategy will be met, and the strategy will not.

Thirty-eight minutes

Go back to the adviser at the second desk, because she is the person the number is most likely to lose. She is not slow. She is the one the others ask when a customer's equipment will not behave, and the one who stays on the line until it does. Every long call she takes to fix a fault properly pushes her own figure up and her team's figure with it, and every Friday the screen tells her, in a colour, that she is part of the problem. Nobody has said so. Nobody needs to.

The canon describes what that does to a person, in a passage from the same book on performance systems, and it is worth reading slowly:

> *When a person's full contribution is systematically reduced to a set of metrics that capture only a fraction of what they bring, the body registers the incongruence. There is a specific somatic signature to this experience — a tightening in the chest, a constriction in the throat, a heaviness in the shoulders. It is the body's way of saying: I am more than this. What I offer is not being received. Something essential about me is being made invisible.*

The organisational consequence follows quietly. The best adviser either learns to transfer the hard calls at the nine-minute mark, like everyone else, or she leaves for a firm that does not count her this way, and takes with her the knowledge of how to fix a router over the phone for a man in his eighties. Either way the firm has lost the thing its objective was asking for, and the screen will record the loss as an improvement.

The canon puts the general rule in a sentence that turns a familiar management saying round. The maxim *what gets measured gets managed*, it says, "contains a shadow that is rarely acknowledged: what does not get measured does not get managed — and often ceases to be valued at all." It adds the consequence for every firm that prints its values on a wall and pays for something else: "if the performance system rewards individual output, speed, and compliance, the posters are wallpaper." The objective at the top of the contact centre's plan is, at this moment, wallpaper.

None of this is an argument for removing the number. A contact centre that does not know how long its calls take cannot staff its shifts, and a firm that cannot see its handling time will not notice when a new system makes every call slower. The argument is narrower and more useful. A number set alone will drift toward what it can see. The task is to set it so that it cannot drift without someone noticing.



The people just short of the line

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER THREE

The people just short of the line

There is one more finding to carry from the research into the contact centre, and it concerns where on the wall the pressure falls. Schweitzer, Ordóñez and Douma found that the people most likely to overstate their performance were not the ones far from the goal but the ones just short of it. Read beside the wall, that is a precise prediction. The advisers whose figures sit a few seconds above the line in the last week of the quarter are under more pressure than anyone else in the room, and the four transfers at the nine-minute mark came from three of them.

The finding suggests that a firm worried about gaming should look not at its worst performers but at the band just below the target, and at the last weeks before a target is assessed. It also suggests something about the design of the target itself. A single line, with a reward or a reprimand on either side of it, concentrates the pressure in exactly the place it does most harm. A range, or a direction of travel, or a learning goal for work nobody yet knows how to do well, spreads it out.

Bevan and Hood's study points the same way at the scale of a national health service. Where consequences for managers were tied to whether a single threshold was met, the threshold itself became the thing that was managed. The lesson they drew was not to abandon targets but to audit them, and to watch for the behaviour that a threshold invites.

None of this requires anyone in the contact centre to be suspected of anything. The advisers just short of the line are, for the most part, doing their honest best under a design that leaves them one move to make. The team leader does not need to watch them more closely. She needs a design in which the move that helps the man in his eighties is also the move that counts.

The pair on the page

Picture the team leader doing it differently in the next quarter. Before the targets are set, her manager asks every team leader to bring one page for each key result, and the page has four lines on it.

The first line is the purpose, in a sentence a new adviser would understand, copied from the objective: customers who never have to call twice. The second is the key result itself, and beside it whether it is committed or aspirational, decided out loud. The third is the paired measure, the number that would rise if the key result were met in the wrong way. For handling time it is easy to find, because the firm already keeps it: the share of customers who call again within seven days about the same fault. The fourth line is a sentence beginning *we will not*: we will not transfer a call to protect the time; we will not end a call before the fault is fixed.

The order of what happens next matters more than the page. The two numbers go on the wall together, the same size, in the same colour, so that nobody can celebrate one without seeing the other. The team leader reviews them as a pair each week, and the first question in the review is not whether the line was met but whether the two moved together. And for the part of the work where nobody yet knows the best method, the long calls with customers who struggle with the equipment, she sets a learning goal in place of an outcome goal, as Seijts and Latham describe: by the end of the month, the team will have found and written down three ways of shortening a router call without leaving the customer to call back. Her most experienced adviser is asked to lead it.

There are failure modes, and they are predictable. The paired measure is chosen because it is easy to improve, not because it would expose the key result. The *we will not* line is written and never read again. The

pair is reviewed, but only the first number is reported upward, so that the second becomes a local matter. And an aspirational result is quietly treated as committed in the year-end review, which teaches everyone to stop setting aspirational ones. Chapter 7 carries the run sheet for this practice, and Chapter 6 places it in the order of a first ninety days.

The gain in this is unusually easy to see, once the pair is on the wall. Every customer who calls twice costs the firm a second call, often a third, and sometimes the customer. Those costs were always in the firm's records, in a different team's report. The pair moves them onto the same wall as the target that was producing them. The difference between the cost of repeat calls before the pair and after it is a gain found, and because both numbers come from records the firm already keeps, it can be measured in a way [both sides could sign]({{BOOK:orgdev-flow}}). The earlier problem in this series, a target that is never allowed to travel back up (Chapter 2), is where the pair should be agreed; and the place it will be read, or not, is the quarterly review that reports in colours (Chapter 4).

Setting the pair in your own building

The article this chapter began as followed one team and one key result. The book can go further, into the finer points of setting paired measures across a set of objectives, and the first of them is finding the second number.

The paired measure is the number that would rise if the first were met in the wrong way. The way to find it is to ask, of each key result, a plain question: *if a sensible, tired person wanted to hit this number without doing what it is for, what would they do?* Then ask what record would show it. For handling time it is repeat calls. For a sales target it is often returns, discounts or cancellations in the following quarter. For a cost target it is often overtime, agency spending or quality faults that appear later and elsewhere. For a delivery date it is rework after the date. In almost every case the second number already exists in the organisation's records, usually in a different team's report. The work is to move it onto the same page as the first.

The second finer point is that the pair is chosen before anyone knows how the quarter will go. A second number added after the first has gone green is an accusation; the same number agreed in advance is a design. The pair, and the line beginning *we will not*, should be written on the same day as the key result and signed by the person who will be held to it.

The third is choosing the kind of goal. Every key result should be marked, out loud, as one of three kinds. A committed result is expected to be met in full, and the team is asked to say early if it will not be. An aspirational result is set high on purpose, and falling short is not a failure. And a learning goal, in the sense Seijts and Latham describe, is set where the team does not yet know how to do the work well: its aim is to find and write down a number of ways that work, and the outcome target comes later. Most failures of the method begin with a learning goal dressed as a committed result. The contact centre's handling time for complex calls was such a case.

The fourth is the review of the pair. The two numbers are reviewed together, every week or every month, and the first question is always whether they moved together. If the first improves and the second worsens, the improvement is not yet a gain, and the team's job is to find out why before anyone celebrates. If both improve, it probably is. If the first worsens and the second improves, the team may have found something that matters more than the target, and the target itself may need to change.

The fifth is the band just short of the line. The research described in this chapter suggests that the pressure to game a target gathers among the people just below it, in the last weeks before it is assessed. A leader can design for this without suspecting anybody. Reviewing progress as a direction of travel rather than a single crossing, avoiding consequences that turn on one threshold, and paying particular attention to the second number in the last weeks of a period all spread the pressure out.

The sixth is how many key results there are. A team with twelve key results cannot hold twelve pairs, and in practice will watch the three that are reviewed most often. Fewer is better: three or four objectives, each with two or three key results, each with its pair. The house's reading of the method in the Luminous 100 warns against cascading it mechanically down the organisation, and the practical form of that warning is to let each team write its own key results against the objective above, through the catchball of Chapter 2, rather than inherit a share of the number.

And the decision that most changes the result is what happens when the pair disagrees for the first time. If the leader treats a falling first number and a rising second as a problem with the second number, the pair is over. If she treats it as the most useful information the quarter has produced, and asks the team what it means, the pair will last. Chapter 6 places the pairing in the second month of a ninety-day plan, after the targets have travelled back up, and Chapter 7 gives the pair page, the question for finding the second number and the choice of goal as working pages.

A quarter past four, the next quarter

It is a quarter past four on a Friday, three months later, and there are two numbers on the screen for each team. The first is the handling time, which is a little above the line. The second is the share of customers who called back within the week, which has fallen for the ninth week running. Nobody claps. The team leader notices that two of her advisers have stopped at the end of the row to look at the second number, which they never did with the first.

Two desks away her most experienced adviser is on a call. It has been running for twenty minutes. On the wall beside her desk, pinned at eye level, is a page with three ways of shortening a router call written on it in her own hand.

The screen is green on one side and green on the other, and this time it means something.

CHAPTER FOUR

The Review That Reported Green

A review that asks whether the plan is on track learns only what it already knew.



The Review That Reported Green

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FOUR

Two o'clock, slide four

It is two o'clock on a Tuesday, and the quarterly strategy review has reached slide four. The room is on the top floor of a firm that makes packaging for food, with a long window that looks over the car park toward the loading bays, and the afternoon light is coming in flat across the table. Slide four is the portfolio: fourteen strategic initiatives in a column, each with an owner's name and a coloured circle beside it. Twelve are green. Two are amber. There is no red, and there has not been a red on slide four for three quarters.

The head of digital is fourth on the list, with the new customer ordering platform, and his circle is green. It is green because last month, with the agreement of the programme office, the milestones were moved. The platform is four months later than the plan the board approved in January, and it is on track against the plan as it now stands. He is sitting with his hands flat on the table and a pen he is not using, and as the programme director thanks everyone for a strong quarter and turns to the two ambers, he feels the small, warm relief of a man who has not been asked, and under it something colder, because he knows the relief is the problem.

He is a composite, drawn from the way these reviews are described in the research below and in the house's own practice, and no firm is meant. Most people who have sat in a strategy review have watched slide four go by.

Ammanuel Santa Anna's canon lists the kinds of truth an organisation finds hardest to say, and the first kind it names is the easiest to describe and the most often withheld: "The project that's behind schedule but being reported as on track."[30] It is not describing a lie. The head of digital's circle is accurate. It is describing an instrument built to answer the wrong question.

This chapter makes one claim. A strategy is lost when its review asks whether the plan is being followed rather than what the plan has learned. The traffic light is a reasonable answer to the first question and no answer at all to the second, and an organisation that reviews its strategy by colour will learn, quarter after quarter, only what it already knew.

What the colour measures

A status colour measures one thing well: the distance between the work and the current version of the plan, as judged by the person responsible for the work. Each part of that sentence matters. The distance is to the current plan, not the original one, so a plan that is rewritten every month can stay green for a year. The judgement is made by the owner, who is also the person who will be asked to explain an amber. And the colour says nothing about why the distance is what it is, or whether the plan itself has turned out to be wrong.

The canon describes the result as a symptom of an organisation whose sensing has gone numb. Listing the signs of compromised feedback, it includes this one: "Feedback mechanisms exist (surveys, retrospectives, review processes) but are ritualistic — the data is collected but nothing changes." Slide four is not a failure of honesty. It is a ritual that has kept its form and lost its purpose. The fourteen circles are collected every quarter, and nothing in the strategy has changed because of them in a year.

The 2015 survey by Donald Sull, Rebecca Homkes and Charles Sull, published in the *Harvard Business Review*, named the belief underneath the ritual as one of its five myths of execution: that execution means sticking to the plan.[31] As the study is reported, a substantial share of managers named difficulty in adapt-

ing to changing circumstances as a cause of failed execution, and the firms that moved resources quickly when conditions changed did better than those that did not. The figures are survey results in a practitioner magazine and are taken here from summaries of the article, because the original was not reachable for this book. The direction of the argument is what matters, and it is not new. A plan is a set of predictions about what will work. A review that asks only whether the plan is being kept treats the predictions as settled.

The canon makes the same distinction in a sentence that could be pinned above slide four: "Execution assumes that the plan is correct and that the task is faithful implementation." The review is built for execution. It has no line on which to record that an assumption in January has turned out to be false.

Why green is safe

There is a reason the colours drift toward green, and it is not that owners are dishonest. It is that the review makes amber costly. An amber circle earns its owner the next twenty minutes of the meeting, in front of the executive team, explaining. A green one earns nothing at all. Over three quarters the room learns, without anyone deciding it, that the best way to protect an initiative is to keep it green and to manage the problems somewhere else.

Santa Anna's canon, in its chapter on accountability, describes the mechanism with unusual directness:

> *Accountability, in its conventional organizational incarnation, is a mechanism of retrospective judgment. It looks backward. It identifies fault. It administers consequences. Its emotional signature is fear — the chronic, low-grade anxiety that pervades organizations where people know that their mistakes will be noticed, named, and penalized.*

It then names the consequence that slide four exhibits: "When people are afraid of being held accountable, they hide their mistakes rather than surfacing them early when they can still be corrected." The head of digital's four months were visible in March as a slipping dependency on a supplier's system. In March they were cheap. By the time a later review is forced to colour the platform amber, they will be expensive, and the conversation will be about him rather than about the supplier.

The house has written about [the meeting after the meeting]({{BOOK:orgdev-flow}}), where the doubt a room will not voice at the table is spoken in the lift afterwards. Slide four is the same pattern written into an instrument. The doubt is not even voiced in the lift. It has been absorbed into a rescheduled milestone, and the colour reports the absorption as success.



The questions an army asks

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FOUR

The questions an army asks

There is a long-standing practice built for the opposite purpose, and its best-known version comes from the United States Army. The after-action review, developed in the army's training centres from the 1970s onward and since adopted in many hospitals, airlines and firms, asks four questions after any significant piece of work: what was supposed to happen; what actually happened; why there was a difference; and what will be done differently next time.[32] It is held soon after the event, by the people who did the work, and its rule is that it is for learning, not for blame. Practitioners who have written about it, among them

Marilyn Darling and Charles Parry in 2001, describe it as a way of linking reflection directly to the next plan.[33] Those accounts are practitioner writing and should be read as descriptions of the practice, not as evidence of its effect.

The evidence for the effect comes from elsewhere. In 2013 Scott Tannenbaum and Christopher Cerasoli published a meta-analysis of debriefs, the general family to which the after-action review belongs, gathering forty-six samples with more than two thousand participants.[34] On average, debriefs improved effectiveness over control groups by about a quarter, an effect size the authors reported as d = .67, and the effect was similar for teams and individuals, in simulated and real settings, in medical and non-medical work. They also found that debriefs worked better when participants, focus and measures were aligned, and when the debrief had structure and facilitation. The studies are mostly of training and of discrete tasks, not of strategies running over a year, and that limit should be held in mind. What the finding supports is narrower and still useful: a structured conversation about what was expected and what happened makes the next attempt better, by a margin large enough to notice.

One study in that family bears directly on slide four, because slide four only ever discusses its ambers. In 2005 Shmuel Ellis and Inbar Davidi ran a field experiment with soldiers on successive navigation exercises.[35] Soldiers debriefed on both their failures and their successes after each day improved significantly more than soldiers who reviewed only their failures. Before the experiment, the soldiers' understanding of their failed events was richer than their understanding of their successful ones; reviewing the successes closed the gap. A green initiative, in other words, is a lesson the organisation has not learned yet, because nobody has asked why it went well, or whether it really did.

The house has written about how [the order of the questions in a weekly review decides what the review teaches]({{BOOK:rsd-rsd-and-self-confidence}}). The quarterly review is the same instrument at the scale of a firm, and its order teaches the same thing.

Check, then act

The quality movement that gave the world hoshin kanri had its own word for what slide four leaves out, and it sits in the middle of the cycle most managers learned early: plan, do, check, act. The cycle is usually traced to Walter Shewhart's work on statistical control at Bell Laboratories in the 1920s and 1930s and to W. Edwards Deming, who carried it to Japan after the war; that lineage is the standard account and is summarised here, not reopened.[36] What matters for the review is the order of the last two words. The check is not a verdict on the people who did the work. It is a comparison of what the plan predicted with what happened, made so that the next action is better than the last.

When Barry Witcher and Rosie Butterworth studied hoshin kanri in three Japanese-owned manufacturing subsidiaries in Britain, they found it was driven by a process of review, and that administering the periodic review was one of its main practical difficulties.[37] The finding is from case studies in three firms, and it is

not a measurement of effect. But it locates the weak point exactly where slide four sits. A firm can deploy its objectives with great care, through every level, and still lose them in a review that only asks whether each level is on schedule.

The canon describes what a strategic initiative looks like when it is held in the spirit of the check rather than the verdict. Writing about governance in a world that cannot be fully predicted, it says that experiments should be designed with intention, with hypotheses, measures, time horizons and criteria for stopping. "But they are held *as experiments* — not as commitments carved in stone." An initiative on slide four is a commitment carved in stone. Its owner can only report that the stone is still standing, or move the stone.

The difference is not academic for the head of digital. If his platform were held as an experiment, its January plan would have said what the firm believed: that customers would order online once the platform existed, that the supplier's system would connect without trouble, and that the sales team could change shape by the winter. Each belief would have had a date by which it could be checked. In March, the second belief would have failed its check, and the review would have had something to learn. Instead it had a colour.

What the review never sees

There is a further cost, and it is the one that most directly separates the strategy from the floor. A review of fourteen initiatives by colour sees the initiatives. It does not see the assumptions that connect them to the strategy, and it does not see the work below them. When the customer ordering platform is late, the strategy assumed that customers would order online by the autumn and that the sales team could be smaller by the winter. Neither assumption appears on slide four. The sales team has already been told.

The canon describes the alternative in its chapter on commitments, as a promise built into every commitment from the start: "*I commit to communicating honestly about my progress, including when I am falling behind, encountering obstacles, or realizing that the commitment needs to be renegotiated.*" It adds the reason that promise is worth more than any colour: "When the commitment to communicate is genuine, problems surface early — when they can still be addressed." A commitment to renegotiate is not a weakness in a plan. It is the plan's way of learning.

The gain here is easy to describe and almost never counted. Every assumption that fails in a strategy will be discovered eventually; the only question is when. Discovered in March, the platform's four months are a conversation with a supplier and a later date for the sales team. Discovered in October, they are a recovery: contractors kept on, customers given apologies, a sales team rebuilt. The difference between the two is a real cost that belongs to no line in any report, and it can be measured, from the firm's own project and cost records, in the way the house has described for [a number both sides can sign]({{BOOK:orgdev-flow}}).



Who chooses the colour

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FOUR

Who chooses the colour

There is a quieter question underneath slide four, and it decides more than any other: who chooses the colour. In most firms the answer is the owner of the initiative, sometimes moderated by the programme office. The owner is the person who will spend the next twenty minutes explaining an amber, and the programme office is measured, in part, by how many of its initiatives are green. Neither is dishonest. Both have reasons to prefer one colour.

Santa Anna's canon describes the general mechanism in a passage on information in hierarchies. Information travelling upward, it says, is at each layer "filtered, summarized, and edited — often unconsciously — to match what the next layer expects or wants to hear." A status colour is the most compressed summary an organisation produces: a whole quarter of work reduced to one of three words. The compression is useful. It also leaves no room for anything the next layer does not expect.

The remedy is not to distrust owners. It is to change who is in the room when the colour is chosen. A colour set by the owner together with the person nearest the work, the developer working on the supplier's interface, the planner who knows the delivery dates, carries information the owner alone cannot supply, and it is harder to round toward green, because two people have to agree on it. It also changes what the colour means to the person nearest the work. Instead of hearing about the strategy from slide four, months later, she has helped to write it.

There is one more rule that matters, and it is about timing. An amber reported early should cost its owner less than an amber reported late, and the firm should be able to show that it does. If the first early amber is met with twenty minutes of questions and the late one with twenty minutes of questions, the room will learn that timing makes no difference, and it will go on reporting late. If the early one is met with thanks and help, and the late one with the question *when did you first know*, the room will learn the difference in a quarter.

The learning review

Picture the next quarterly review, and a programme director who has decided to change one thing. She does not abolish slide four. She moves it. The colours go out in writing three days before the meeting, set not by each owner alone but by the owner together with the person nearest the work, and every owner is told, in writing and in advance, that an early amber will be thanked and never questioned on the day it appears.

The meeting itself is spent on three initiatives, not fourteen. She chooses them by what the room can learn, not by colour: one amber, one green that has moved its milestones, and one green that seems to be going well. For each, the owner answers the four questions of the after-action review in ten minutes: what the strategy expected by now, what has actually happened, why there is a difference, and what the owner would change. The head of digital, asked about his green platform, says that the strategy expected online ordering by the autumn and will not get it until the spring, and that the reason is a supplier's system nobody had tested, and that the sales team's plan should change. The room is quiet for a moment. Then the chief commercial officer asks the next question about the supplier, and not about him.

The last fifteen minutes are kept for one thing: what in the strategy itself changes because of what the room has just heard. Not the initiatives; the assumptions. The programme director writes them on a single page, each with an owner and a date, and the page goes to the board with the colours.

The canon describes reviews conducted in this spirit in a sentence that is also a standard: "These retrospectives, conducted with genuine curiosity and without self-congratulation, are among the most powerful developmental practices available to governance groups." The test of the standard is simple. After the review, has anything in the strategy changed, other than a date?

There are failure modes, and they are predictable. The programme director thanks the first early amber and then, in the next quarter, asks its owner why the problem was not foreseen. The three initiatives are always the troubled ones, so that being chosen becomes a sign of trouble. The page of changed assumptions is written and never read by the board. And the four questions are answered in the language of the slides, smoothly, because nobody has yet seen what happens to a candid answer. Chapter 7 carries the run sheet for this practice, and Chapter 6 places it in the order of a first ninety days.

The review is where the strategy's numbers are read, which is why a number set without its pair (Chapter 3) will mislead it, and it is also where the firm has its best chance to notice that the strategy it is actually following is not the one it wrote (Chapter 5).

Running the learning review

The article this chapter began as followed one quarterly review. The book can go further, into the finer points of changing an organisation's review rhythm, and the first of them is where to begin, because a quarterly review happens only four times a year and a first attempt should not wait three months.

The learning review is best begun in a monthly operating review, or in one division's review, before it reaches the quarterly strategy review. A monthly meeting gives three attempts in a quarter and a chance to make the mistakes in a smaller room. When the quarterly review changes, the people in it have already seen the four questions work, and the first candid answer at the top table is less of a risk because it is not the first anywhere.

The second finer point is the pre-read. The colours, the figures and the written updates go out at least three days before the meeting, and the meeting does not repeat them. A meeting that walks through the status slides again has spent its time on the part of the review that could have been read. The rule is that anyone who has not read the pre-read does not ask questions about it in the room; the room is for what cannot be read.

The third is choosing what to review. Three items are enough for an hour, chosen by what the room can learn, not by colour alone: usually one amber, one green that has moved its milestones or its scope, and one green that appears to be going well. The third is the one most often skipped and, on the evidence Ellis and Davidi gathered, one of the most valuable, because a success that nobody has examined is a lesson the organisation has not yet learned. The items should rotate, so that being chosen is ordinary rather than a sign of trouble.

The fourth is the four questions and what each is for. *What did the strategy expect by now?* is answered from the original plan, not the current one, because the point is to see what has been learned since January. *What has actually happened?* is answered in facts and dates. *Why is there a difference?* is where the time goes, and its answer should name an assumption, not a person. *What would you change?* asks the owner, first, before the room offers anything. Ten to fifteen minutes an item is usually enough. The chair's job is to keep the third question about the assumption and to stop the room from turning it into the question *who*.

The fifth is the colour rule. The colours are set by the owner together with the person nearest the work, and an early amber is thanked and never questioned on the day it appears. A late amber, one that appears after it could have been seen, is met with a single question asked without heat: *when did you first know?* Over two or three cycles the difference between those two receptions teaches the room more than any statement about candour could.

The sixth is the page that leaves the room. The last fifteen minutes of every review are kept for one thing: what in the strategy itself changes because of what was heard. Each change is written on a single page with an owner and a date, and the page goes to whoever owns the strategy, usually the board or the executive team. If the page is empty for two reviews running, either the strategy is perfect or the review is not yet real, and it is almost never the first.

And the decision that most changes the result is what happens to the first owner who answers the third question honestly about her own initiative. If the room moves to the assumption, the supplier's system, the forecast that was wrong, and the owner leaves the meeting with help, the practice will hold. If the room moves to the owner, the next review will be held in the old language, smoothly, and the four questions will become a ritual of their own. Chapter 6 places the learning review in the second half of a ninety-day plan, beginning in a monthly meeting, and Chapter 7 gives the pre-read rule, the choice of three, the four questions, the colour rule and the page of changed assumptions as a run sheet.

Two o'clock, the next quarter

It is two o'clock on a Tuesday, and the quarterly review has reached the part of the afternoon where slide four used to be. The colours went out on Friday and nobody has mentioned them. The light across the table is the same flat light from the long window, and the loading bays below are busy.

The head of digital is speaking about the supplier's system, which has now been tested, and about the spring. He has a pen in his hand this time, and he is writing on the page of changed assumptions as he talks.

On the screen behind him there is no slide at all, only the four questions and a blank space under the last one.

CHAPTER FIVE

The Strategy You Already Follow

Every organisation follows two strategies, the one it wrote and the one its decisions add up to.



The Strategy You Already Follow

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FIVE

Nine o'clock, the kitchen table

It is nine o'clock on a Sunday evening in January, and the chief executive of a firm that makes and services commercial kitchen equipment is at her own kitchen table with a laptop, a glass of water and a spreadsheet her finance director sent on Friday. It lists every request for money, people and time that the executive team approved last year, one line each, with the date and the reason. There are three hundred and twelve lines. The house is quiet. Above the table a single pendant lamp makes a circle of warm light on the wood, and the rest of the room is dark.

The strategy the board approved a year ago fits on one slide, and she knows it by heart: grow the service business, so that by the end of the decade the firm earns more from keeping kitchens running than from selling them. She sorts the spreadsheet by purpose, twice, to be sure. Most of last year's money went to the equipment business: a new production line, a larger sales team for the new range, a second warehouse. Each request was sensible. Each was approved on its merits, in a meeting she chaired. The service business received what it asked for, which was not much, because the service managers had stopped asking. Her stomach does something slow and unpleasant. Nobody decided to abandon the strategy. She is reading the record of an organisation that followed another one, one reasonable decision at a time.

She is a composite, drawn from the way these evenings are described in the research below and in the house's own practice, and no firm is meant. Anyone who has read a year of their own approvals has met the second strategy.

Ammanuel Santa Anna's canon opens its chapter on emergent strategy with a story of a plan meeting the world,[38] told as the canon's own account of a healthcare cooperative's annual retreat:

> *In the summer of 2019, the leadership team of a mid-sized healthcare cooperative in the American Midwest gathered for their annual strategic retreat. They spent three days in a conference center, surrounded by whiteboards covered in SWOT analyses, competitive positioning maps, and cascading five-year goals. By Sunday evening, they had produced a strategic plan of impressive comprehensiveness — eighty-seven pages of objectives, initiatives, timelines, and key performance indicators. The document was professional, internally consistent, and unanimous in its approval.*

Six months later, the canon goes on, the world changed and the plan became, in its words, "a memorial to a world that no longer existed." The chief executive's plan has not been overtaken by a pandemic. It has been overtaken by her own organisation.

This chapter makes one claim. Every organisation has two strategies: the one it wrote, and the one its decisions add up to. A strategy is lost on the way to the floor when nobody compares the two. It is sometimes found there too, because the second strategy is occasionally wiser than the first, and the only way to know which is to read it.

A pattern in a stream of decisions

In 1978 Henry Mintzberg published a paper in *Management Science* that changed what the word strategy could mean.[39] The usual definition, he observed, encouraged the idea that strategies are deliberate plans conceived in advance of the decisions that carry them out. He proposed another: a strategy is "a pattern in a stream of decisions." The definition made it possible to study both the strategies organisations intended

and those they realised despite their intentions. His two examples were long ones, the strategies of Volkswagenwerk from 1934 to 1974 and of the United States government in Vietnam from 1950 to 1973, and his conclusion was that the interplay between intended and realised strategy might lead to the heart of how organisations actually choose.

In 1985, with James Waters, he set the idea out as a continuum.[40] At one end lies the purely deliberate strategy, intended precisely and realised exactly as intended; at the other the purely emergent one, a consistent pattern that nobody intended. Real strategies lie between, and the two authors named the kinds they had found in their research, among them planned, entrepreneurial, umbrella, process and imposed. An umbrella strategy sets broad guidelines and lets the specifics emerge beneath them. An imposed one is set from outside, by a parent company or a government.

Mintzberg's work is built on long historical case studies, read closely, and its strength is description rather than measurement. It does not say that emergent strategy is better than deliberate strategy, and Mintzberg was careful not to. It says that both are always present, and that a leadership team which looks only at the deliberate one is looking at half of what it is doing.

Santa Anna's canon, which acknowledges Mintzberg by name, puts the idea in a sentence that the chief executive could write at the top of her spreadsheet: "Strategy, in this view, is not what you plan. It is what you *do*." Three hundred and twelve approvals are what her firm did.

The rule that decided

The most closely studied case of a realised strategy overtaking a written one comes from Intel, and it was documented by Robert Burgelman of Stanford over many years of fieldwork inside the company.[41] In the early 1980s Intel thought of itself as a memory company. Memory chips were its founding business and the centre of its identity. Yet by the middle of the decade most of its manufacturing capacity was being used for microprocessors, and Intel's formal decision to leave the memory business, taken by its top management in 1985, ratified a shift that had already largely happened below them.[42]

The mechanism, as Burgelman described it, was not a rebellion. It was a rule. Intel's middle managers allocated scarce factory capacity by a simple criterion, the margin earned on each start of a silicon wafer, and as memory chips became a commodity and their margins fell, the rule moved capacity, quarter by quarter, toward microprocessors, whose margins were higher. Nobody in the factories decided that Intel should become a microprocessor company. The allocation rule decided it, and the strategy on paper followed some years later. Andrew Grove later told the story of the decision from the top in his book *Only the Paranoid Survive* (1996).[43]

The account here rests on Burgelman's published work and on summaries of it; the 1994 paper's full text was not reopened for this book, and some details of dates differ between secondary accounts. What the

case shows, and what Burgelman drew from it in his 1991 paper on the ecology of strategy making inside firms, is that strategy is shaped by two processes at once: an induced process, in which the formal strategy guides decisions, and an autonomous one, in which people below the top respond to what they see and the organisation's own rules select among their responses. Successful reorientations, he proposed, are likely to have been preceded by internal experiment and selection of exactly this kind.

The chief executive's firm has an allocation rule too, though nobody has written it down. The rule is that a request with a clear return in the next two years is approved and a request whose return lies further out is deferred. It is a sensible rule. It is also, in the arithmetic of three hundred and twelve lines, a strategy, and it is not the one on the slide.



Read in the allocations

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FIVE

Read in the allocations

If the realised strategy is a pattern in a stream of decisions, the practical question is where the stream can be seen. The answer, for most organisations, is in the places where scarce things are given out. Joseph Bower's study of how large firms allocate capital, published as *Managing the Resource Allocation Process* in 1970, argued that strategy in a diversified firm is made less in the planning office than in the long chain of proposals, sponsorships and approvals through which investments are actually chosen, often starting well below the top.[44] Bower's book is a field study of a small number of firms, summarised here from its stand-

ard reading rather than reopened. Its idea has proved durable because anyone can test it against their own records. A firm's strategy is visible in its approvals, its hiring and its senior people's diaries, whether or not anybody has looked.

The three are not equal witnesses. Money is the easiest to count and the most often studied. Hiring is slower to move and harder to reverse, so it shows what the organisation expects to be doing in two years. And the diaries of the executive team are the most revealing of all, because senior attention is the scarcest thing a firm has and the one nobody approves. The chief executive's calendar for last year would show, if she opened it, that she spent four afternoons with the service managers and most Thursdays with the equipment sales team. Nobody decided that either.

Santa Anna's canon describes the organisation that grows this way, without a plan, as an image from the forest. "A forest does not have a strategic plan," it says. "It has a pattern of adaptive response — millions of organisms, each responding to local conditions, collectively producing an order that no central planner could have designed." The image is not an argument against planning. It is a description of what is always happening underneath a plan, in any living system, and of why the pattern deserves to be read as information rather than dismissed as noise.

The house's Luminous 100 adds a practical note about the kind of problem the kitchen table presents. Of the Cynefin framework, Dave Snowden's sorting of decisions by how far cause and effect can be known in advance, it says that it is valuable "because it tells you WHAT KIND of problem you're facing before you try to solve it." Whether customers will buy service contracts at scale is not a problem that more analysis will settle. It is the kind that is settled by trying something small and watching what happens.

Two strategies, and which is right

There are two ways to read the kitchen table, and the canon holds both of them.

The first reading is that the organisation has drifted, and the written strategy should be enforced. The canon is clear that direction matters: "A governance body that provides no direction is not practicing emergence — it is practicing abdication." A firm that approves every sensible request on its merits, with no view of what the requests add up to, has no strategy at all, only a habit.

The second reading is that the organisation knows something the slide does not. Perhaps the equipment business is growing because customers want the new range more than anyone forecast, and the service managers stopped asking for money because their customers are not yet ready to buy service contracts. Intel's allocation rule was right, and its formal strategy was the thing that was late. The canon describes this possibility in its chapter on living systems, writing of the patterns that arise inside every organisation from the relationships between its people, "patterns that no manager designed and no strategic plan anticip-

ated," and arguing that they are often the system's own intelligence rather than inefficiencies to be removed.

It also offers a way of holding the two readings together, which is the most practical thing in the chapter. Purpose, it says, is to be held firmly, strategy lightly and tactics very lightly. It puts the point as an image:

> *The metaphor is sailing. The destination (purpose) is fixed. The general course (strategy) adjusts to wind and current. The moment-to-moment handling of the sails (tactics) responds to the immediate conditions. A sailor who fixes the sails in one position and refuses to adjust them regardless of changing wind is not a committed sailor. That sailor is a drowning one.*

And it names the danger on the other side, the firm that follows its written strategy faithfully into a world that has changed: "the more faithfully the organization executes an outdated plan, the further it moves from the reality it needs to serve."

Which reading is right at the kitchen table cannot be known from the spreadsheet alone. That is the point. It can only be known by putting the two strategies side by side, on purpose, and asking.

The rule nobody wrote

The most useful thing the chief executive can find at the kitchen table is not a decision. It is a rule, because rules are what make three hundred and twelve decisions point the same way. Intel's rule was written down: margin per wafer start. Hers is not, which makes it both more influential and harder to see.

There is a simple way to find it, and it needs nothing but the spreadsheet and an hour of the executive team's time. Take the last twenty approvals and the last twenty refusals or deferrals, and ask, of each, the reason it was given at the time. Not the reason on the form; the reason people remember. Within an hour the same few reasons will appear again and again: *it pays back inside two years*, *the customer asked for it*, *we already have the people*, *it is what we know how to sell*. Those sentences are the firm's realised strategy, written in the voice of its approvals. They will rarely mention the slide.

Once the rule is written, it can be judged. Some unwritten rules are wise, and the written strategy should learn from them. Others are habits that once made sense and now quietly overrule the board. The two-year payback rule at the kitchen-equipment firm is both. It protected the firm through a hard decade, and it is the reason the service business, whose returns arrive slowly and over many years, has never been funded.

The canon has a sentence on why the effect of a strategy is so hard to predict from its size, and it applies with particular force here. In a living system, it says, "Large inputs can produce negligible effects (the multimillion-dollar change initiative that changes nothing; the comprehensive strategic plan that gathers dust)." A plan the size of a slide deck can be outweighed by a rule the size of a sentence, if the rule is the

one applied every Thursday. That is not a reason to despair of plans. It is a reason to write the rule down, and to put it on the same page as the plan.



What the comparison is worth

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER FIVE

What the comparison is worth

Before the practice, the gain, because here it is largest and least visible. When an organisation's decisions add up to one strategy and its words to another, it pays twice. It pays for the investments the decisions make, and it pays for the plans, people and promises built on the words: the service managers hired for a

business that is not being funded, the board told a story that the money does not tell, the customers promised contracts that nobody is building. And if the written strategy is the wiser one, it pays a third time, in the value of the service business that the allocation rule quietly declined to grow.

None of these costs appears in any report as the price of two strategies. They appear as missed targets in one division and surprise growth in another, each explained separately. The difference between what the firm earns when its decisions and its strategy agree and what it earns when they do not is a real gain, sometimes a very large one, and it can be found and measured from records the firm already keeps: the approvals, the hires, the calendars, the results. The house has written about how such a gain should be [measured before anyone shares in it, honestly enough that both sides would sign the number]({{BOOK:orgdevflow}}). This book does not describe how a measured gain is divided; that depends on the organisation and on the people whose work produced it.

The quarterly reading

Picture the chief executive doing one thing differently. She does not call a strategy offsite. In the first week of every quarter she asks her finance director for three lists, all drawn from records the firm already keeps: where last quarter's money went, sorted by the strategy's own priorities; where the new hires went; and, from the executive team's calendars, where the senior hours went. She lays them out, on one page, beside the slide.

The executive team spends one meeting a quarter on that page, and only on the places where the two strategies disagree. For each divergence it gives one of three answers, out loud and on the record. The first is *adopt*: the pattern is wiser than the plan, and the written strategy changes to say so. The second is *stop*: the pattern is drift, and the rule that produced it changes. In her firm that means writing down the unwritten rule about two-year returns and deciding, explicitly, how service requests will be judged instead. The third is *test*: nobody knows which strategy is right, so the team designs a small, contained experiment to find out, with a hypothesis, a measure, a date and a point at which it will stop. The canon describes this third answer in its own terms, as the safe-to-fail experiment, designed so that failure teaches rather than harms.

The first quarter is uncomfortable. The page shows that the service business received a small share of the money and almost none of the senior hours, and the executive team learns that it has been approving one strategy and describing another to the board. The team chooses *test*: three service contracts sold at cost to existing customers, to find out whether the demand the strategy assumed exists. Two quarters later the answer is in, and the strategy is rewritten, a little, to match what was learned.

There are failure modes, and they are predictable. The page is assembled but the meeting is spent defending the divergences rather than deciding them. Every divergence is answered *stop*, because the team cannot bear to admit the plan was wrong, or *adopt*, because it cannot bear to admit the decisions were. The calen-

dars are left out because they are the most personal list. And the allocation rule, once named, is not changed, because changing it means saying no to sensible requests. Chapter 7 carries the run sheet for this practice, and Chapter 6 places it in the order of a first ninety days.

The reading is where the rest of this book meets. The three good words on a card (Chapter 1) become a strategy only when they decide something; the review that reports in colours (Chapter 4) becomes useful only when it asks which assumptions have failed. The quarterly reading asks the same question of the whole firm at once.

Reading your own realised strategy

The article this chapter began as followed one chief executive and one spreadsheet. The book can go further, into the finer points of reading the realised strategy every quarter, and the first of them is assembling the three lists honestly.

The money list is usually the easiest, because finance already keeps it: every approval of spending above a level the executive team chooses, with its date, its sponsor and the reason given. The hiring list comes from the people records: every new role created or filled, and the part of the organisation it went to. The hours list is the hardest and the most revealing, and it should come from the executive team's own calendars for the quarter, sorted roughly by the strategy's priorities. It is personal, and some teams will want to leave it out. It should not be left out, because senior attention is the one resource no approval process records, and it is often where the second strategy is most visible.

The second finer point is sorting by the strategy's own words. Each line on each list is placed under the priority it serves, or under *none of these*. The sorting should be done by two people who did not make the decisions, usually someone from finance and someone from the strategy or planning team, and their disagreements noted. The *none of these* column is not a failure. It is often where the firm's real business lives, and it is the first place to look for the realised strategy.

The third is finding the unwritten rule. Take the last twenty approvals and the last twenty refusals or deferrals, and ask the people who decided them for the reason as they remember it. Write the reasons down in their own words and count them. Within an hour the same few sentences will appear repeatedly. Those sentences are the allocation rule the organisation is actually using, and they should be written at the top of the page, beside the strategy, in the same type size.

The fourth is the three answers, and the discipline of using all three. For each place where the lists and the strategy disagree, the executive team says *adopt*, *stop* or *test*, out loud and on the record. A team that only ever says *stop* is defending its plan against its own organisation. A team that only ever says *adopt* is letting the organisation's habits write its strategy. The proportion is not fixed, but over a year a healthy team will

use all three, and will be able to point to at least one place where its written strategy changed because its realised one was wiser.

The fifth is the test. A *test* is a small, contained experiment designed to find out which of two strategies is right, in the spirit of the safe-to-fail experiment the canon describes: a hypothesis, a measure, a date, and a point at which it stops. It should be small enough that failure is affordable and specific enough that the result will actually decide something. Three service contracts sold at cost to existing customers is a test. A pilot of the whole service business in one region for two years is a programme, and will be defended rather than read.

The sixth is the rule itself. Once the unwritten allocation rule is on the page, the executive team can decide whether to keep it, change it, or write a second rule beside it for the part of the business it is starving. Changing a rule usually goes further than changing a hundred decisions, because the rule makes the next hundred. It is also harder, because the rule usually exists for a good reason, and the people who apply it every Thursday have been right to.

And the decision that most changes the result is whether the first quarterly reading is allowed to embarrass anyone. It will show that the executive team has approved one strategy and described another. If the team treats that as a finding about the organisation's rules, it will read the next quarter with interest. If it treats it as a finding about who approved what, the lists will be quietly tidied before the next reading, and the second strategy will go back underground.

What the reading finds is also, in most firms, the largest gain this book points to: the value released when an organisation's decisions and its strategy are made to agree, in whichever direction. It can be measured from the same records the lists come from, against a history fixed before the reading begins, in the way [Where the Flow Stops]({{BOOK:orgdev-flow}}) describes in its chapter on a number both sides can sign. This book does not describe how a measured gain is divided; that depends on the organisation, on its people and on its history, and it is decided separately with those who will be covered. Chapter 6 places the first reading in the last month of a ninety-day plan, and Chapter 7 gives the three lists, the sorting, the search for the unwritten rule and the three answers as working pages.

Nine o'clock, a year later

It is nine o'clock on a Sunday evening in January, a year later, and the chief executive is at the same table under the same lamp. The spreadsheet from her finance director has four columns now, not one: the money, the hires, the senior hours, and a fourth column headed *adopt, stop or test*, filled in, quarter by quarter, in the executive team's own words.

She sorts it by purpose once. The service business has more of the money than it had, and more of the hours, and a written rule for how its requests are judged. The equipment business has less than it asked for and more than the slide once allowed, because one of last spring's tests showed the new range was worth it.

The two strategies are not the same yet. For the first time, she can see both of them.

A NOTE FROM THE HOUSE

Every organisation follows two strategies, the one it wrote and the one its decisions add up to. The distance between them is where a gain is usually kept. Luminous Prosperity Inc. reads that distance inside organisations as a gainshare, and is paid as a share of the gain it helps find.

REQUEST A CONVERSATION

CHAPTER SIX

The First Ninety Days

Five practices set in order, from reading the strategy as the floor reads it to reading the strategy the firm already follows.



The First Ninety Days

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SIX

Monday, day one

It is twenty past seven on the first Monday in February, and the chief operating officer of a firm of about six hundred people is at her desk before anyone else on the floor. The strategy was approved by the board in December. It fits on one slide and three cards, and it has been launched: a video from the chief executive, an all-staff meeting, the cards on every desk. She has read the first five chapters of a book like this one, or lived them, and she has a suspicion that the strategy has been distributed and not yet deployed. She has

a notebook, a cup of coffee and the diary of every standing meeting she chairs. She does not have a plan, and she feels in her jaw the familiar pull to write one before nine and present it on Thursday.

She is a composite, drawn from the way these first weeks are described in the research the earlier chapters read and in the house's own practice, and no firm is meant. This chapter is written to her, and to anyone in her position: the person responsible for making a strategy real who wants to begin. It describes the first ninety days in order, with who is in the room, what each step needs, roughly how long it takes, the decisions that change the result, the mistakes that are made at each step and how to see them early, how to read what comes back, and what to do on the first Monday. The working pages it refers to are gathered in Chapter 7.

Two things should be said before the plan. The first is that nothing here should be launched. A strategy that has just been launched does not need a second launch about its launch, and the canon's warning about governance applies with force:[45] "Adding process to an unconscious governance system is like adding lanes to a highway that is going in the wrong direction." The practices in this book are not a programme. They are changes to five things the organisation already does: how it explains a priority, how it sets a target, how it measures progress, how it reviews, and how it allocates. The second is that the order matters more than any single step. The chapters were written in the order a strategy travels down the building. The ninety days run in the order the practices depend on each other.

Why this order

The five practices are the back-brief (Chapter 1), catchball (Chapter 2), the pair (Chapter 3), the learning review (Chapter 4) and the quarterly reading of the realised strategy (Chapter 5). In the first ninety days they run in nearly that sequence, preceded by ten days of reading.

The reasoning is this. Every practice asks people to send something upward that they have been keeping to themselves: what the strategy actually asks of them, whether a target can be met, what a number is hiding, why an initiative is late, where the money really went. People send such things when they have seen that what is sent is answered. So the first change is made where the leader controls the conditions and everyone can watch the result: in the meaning of one priority, settled at the top and tested through a back-brief. Only once a back-brief has been answered does it make sense to throw a target down and ask for it back, because catchball depends on the same trust and asks for more of it. The pair is set on the targets that came back, because a second number is easiest to agree when the first has just been agreed honestly. The learning review reads those pairs, and so it comes after them. The reading of the realised strategy comes last, because it is the most exposing for the people at the top, and they should have seen the other practices work first.

One thing is done out of order, on purpose. The history against which any change will later be read is copied and dated in the first ten days, before any practice could move it. The canon puts the reason for be-

ginning with perception in a sentence: "The discipline of sensing before strategizing does not mean never strategizing. It means ensuring that strategy is grounded in accurate perception rather than wishful projection." The first ten days are spent perceiving.

Who is in the room

Three roles carry the first ninety days, and they should be named before anything else is done.

The sponsor is the most senior person whose decisions the practices will touch: often the chief operating officer, sometimes the chief executive, in a large firm the head of a division. The sponsor's job is to go first, in her own meetings and her own back-brief, to answer what comes up quickly, and to protect the people who send bad news. Expect it to take four to six hours a week, most of it inside meetings she already holds.

The working group is small, five to seven people. It should include someone from planning or strategy, someone from finance who will later read the history and the realised strategy, two managers from the middle of the organisation who carry targets down and results up, one person within a level of the work itself, and, deliberately, one person who joined recently, because newcomers still notice when a priority means nothing to the people who repeat it. The group meets for an hour a week. It reads what the practices produce and brings back what it finds. It should not include anyone whose presence would make the middle managers cautious, and the sponsor should ask them privately, not in the group, whether that is so.

The convenor keeps the record: the floor readings, the back-briefs and their answers, the returned balls, the pairs, the review pages, the three lists, and the dates of each. Often this is the person from planning. The convenor needs about three hours a week, and the quality that matters most is that people trust the convenor to write down what they said rather than what the sponsor would like to hear.

One decision changes the result more than any other at this stage. The sponsor must decide whether she will write her own back-brief, upward, about the strategy as she received it from the board, and show it to the working group. A sponsor who asks the middle for candour about the strategy and offers none of her own has told everyone what the practice is for.

Days one to ten: reading before changing

The first ten days change nothing visible. They gather four inputs.

The first is the floor reading. The sponsor, or a member of the working group whom people trust, asks ten people the question of Chapter 1, privately and without notice: what does the strategy ask you to do differently on Monday? At least half are within two levels of the work. The answers are written down word for word and sorted into three kinds: the card repeated, a manager's instruction repeated, a real choice described in the person's own words. Chapter 7 gives the page.

The second is the truce test, held with the leadership team in one meeting of about ninety minutes. For each priority on the card, the convenor brings one ordinary case from last month, drawn from the floor reading where possible, and each executive writes, privately, what the priority means for that case. The answers are read aloud together. Where they differ, the difference is the first finding of the ninety days. It is not resolved in that meeting; it is written down.

The third is a list of the organisation's targets for the year and of the reviews that read them: which meetings review which numbers, how often, and by colour or otherwise. The convenor notes, for each key target, whether a second number already exists in the organisation's records that would rise if the first were met the wrong way. Nothing is changed yet.

The fourth is the history. Finance and the convenor copy, date and file, unaltered, the records that the later practices will be read against: the last two years of targets and results for the priority chosen, the approvals of money and people for the last four quarters, the review packs, and whatever second numbers exist. This is the history the measurement conversation will read on day sixty and the realised-strategy reading on day seventy-five, fixed before any practice could move it.

At the end of the ten days the working group meets and writes, on a single page, its first guess: which priority is most a truce, which targets will not hold, which reviews report green over a known problem, where the money seems to go. It is a guess, labelled as one, and it is kept for day ninety.

What goes wrong at this stage. The commonest mistake is to announce. A sponsor who tells the organisation that a strategy deployment programme has begun will receive back-briefs written for a programme. The early sign is a question in a corridor: *what is this new thing called?* The second mistake is to begin with a survey of strategic alignment, which tells people they are being measured before anything has changed and produces a score that describes what they were willing to tick. The third is to try to resolve the truce in the first meeting it appears. The leadership team will resolve it with a fourth phrase.



Days eleven to thirty: one priority, settled and briefed back

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SIX

Days eleven to thirty: one priority, settled and briefed back

The second and third weeks do two things in order.

First, the leadership team settles the truce on one priority, the one where the truce test showed the widest difference. It does so in writing, with the case from day ten in front of it, and the result is two or three sentences that say which good thing governs when two conflict, and how much risk the firm will carry to honour it. The canon's distinction between what is held firmly and what is held lightly is useful here: "Purpose

is held firmly." The priority's purpose does not move. What the leadership team writes is the current answer to the conflict, which can be revised as the practices produce evidence.

Then the back-brief runs for that priority in one division. Each manager two levels below the leadership team writes a page in her own words: what the priority now asks of her team, the one or two decisions it changes as if-then sentences with a real case, what the team will stop, and what she needs from above. The sponsor's own back-brief goes first. The person who receives the pages reads all of them before replying to any, sorts them by the choice each makes, takes the disagreements upward unedited, and answers every page within the week.

The decision that changes the result is what happens to the first back-brief that says the priority cannot be done as written. If it is answered with a decision and the meaning of the priority visibly changes, the practice will hold. If it is answered with a correction, the next round of back-briefs will be copies of the card.

What goes wrong. The back-briefs are graded for their wording. The pages ask for budget and become negotiations. The person receiving them answers the easy ones and leaves the awkward. Or the leadership team, shown that its settled sentences have been read four ways, rewrites the sentences without asking why. The early sign is easy to read: if the back-briefs are more alike than the floor reading was, people are writing what they think is wanted.

How to read what comes back. By day thirty the convenor's record should show a settled priority, a set of back-briefs, a list of what each team will stop, and a list of what was asked of the level above, with the answers. The list of things to stop is the most useful single output. It shows what the strategy has actually displaced.

Days twenty-one to forty-five: the target thrown back

Overlapping with the back-brief, the working group chooses the targets that serve the settled priority and runs catchball on them, in the division that wrote the back-briefs. If the year's targets have already been announced to the board, catchball runs on the next quarter's share of them, or on the targets for one region, so that a returned ball can still change something.

Each target goes down with the part of the strategy it serves, in a sentence, and with what the level above believes about how it could be reached. Each level below answers the three questions of Chapter 2: what it would take with what is there now; what could be reached if one named thing changed; and what reaching it would cost elsewhere. The returns separate what is committed, what depends on a decision above, and what cannot be reached without a cost to something else the strategy values. The middle managers meet for an hour without their director before returning their balls, to compare what they have heard. The director reads everything before answering anything, and answers in writing within the week.

The decision that changes the result is what the director does with the first returned ball that says a target is wrong. A decision, in writing, that changes a target or a priority, keeps the practice. A restatement of the target ends it.

What goes wrong. The ball is thrown without its reasons, so it can only come back as yes or no. The returns are softened into a single number. The director treats the returns as a negotiation over a shortfall. The early sign is in the third question: if no return names a cost elsewhere, people have not yet believed that the third question is meant.

How to read what comes back. By day forty-five the record should show a set of targets that have been returned and answered, a list of the decisions the level above had to make to keep them, and a list of the costs elsewhere that were named. The last list is where the strategy's real trade-offs are, and it is often where the first gain is.

Days thirty-one to sixty: the pair

Once the targets have come back, each key result for the settled priority is given its pair, at the level that will be held to it. The team writing it answers the question of Chapter 3: if a sensible, tired person wanted to hit this number without doing what it is for, what would they do, and what record would show it? The pair, the purpose sentence and a line beginning *we will not* are written on one page and signed on the same day as the target. Each key result is marked, out loud, as committed, aspirational or a learning goal. Where the team does not yet know how to do the work well, the outcome target is replaced for this quarter by a learning goal.

The two numbers go on the same page, the same size, wherever the first number was shown before. The team reviews them together each week, and the first question is always whether they moved together.

The decision that changes the result is how the leader responds the first time the pair disagrees, the first number improving and the second getting worse. If she treats it as the most useful information the month has produced, the pair lasts. If she treats it as a problem with the second number, the pair is over.

What goes wrong. The second number is chosen because it is easy to improve. The *we will not* line is never read again. Only the first number is reported upward. The early sign is on the page itself: if the second number has not been looked at in a review by week three, it has become decoration.

How to read what comes back. By day sixty the record should show, for each key result, its pair, its kind, and at least four weeks of both numbers. Where they have moved together, the improvement is probably real. Where they have separated, the separation is the finding.

Days forty-one to seventy-five: the learning review

From the sixth week, one monthly operating review that reads the settled priority's numbers changes its form. A monthly meeting is chosen so that the review can be held three times inside the ninety days; the quarterly strategy review changes later, when its members have seen the form work in a smaller room.

The colours and figures go out three days before the meeting, set by each owner together with the person nearest the work, and the meeting does not repeat them. Every owner is told, in writing and in advance, that an early amber will be thanked and never questioned on the day it appears. The meeting reviews three items, chosen by what the room can learn: usually one amber, one green that has moved its milestones, and one green that seems to be going well. For each, the owner answers the four questions of Chapter 4 in ten to fifteen minutes. The last fifteen minutes produce the page of what in the strategy changes because of what was heard, with an owner and a date for each line, and the page goes to the leadership team.

The canon names the reason this is harder than it looks: "This flexibility is harder than it sounds. Organizational culture — and the human psyche — are deeply invested in the appearance of certainty." The first review will be smooth. The second may not be. The third usually produces the first change to an assumption.

The decision that changes the result is what the room does with the first owner who answers the third question honestly about her own initiative. If the room moves to the assumption and she leaves with help, the review holds. If it moves to her, the next review will be held in the old language.

What goes wrong. The pre-read is not read and the meeting walks through it. The three items are always the troubled ones. The page of changed assumptions is empty, or written and never read by the leadership team. A late amber is met exactly as an early one, so that timing teaches nothing.

How to read what comes back. By day seventy-five the record should show three reviews, three pages of changed assumptions, and at least one change to the strategy's own sentences, not only to dates. If there is none, the review is not yet real.



Days sixty-one to ninety: reading the strategy you already follow
HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SIX

Days sixty-one to ninety: reading the strategy you already follow

In the last month the working group assembles the three lists of Chapter 5 for the previous quarter, from the records fixed on day ten and from the quarter since: where the money went, where the new hires went, and where the leadership team's hours went, each sorted by the strategy's own priorities, with a column for *none of these*. Two people who did not make the decisions do the sorting. The working group also finds the unwritten allocation rule, by asking the people who decided the last twenty approvals and twenty refusals for the reason as they remember it, and writes the rule beside the strategy.

The leadership team spends one meeting of about two hours on the page, and only on the places where the two strategies disagree. Each divergence receives one of three answers on the record: adopt, stop or test. A test is designed as Chapter 5 describes, with a hypothesis, a measure, a date and a point at which it stops.

In the same weeks, the working group chooses one gain the earlier practices have found, a cost elsewhere named in a returned ball, a second number that has fallen since the pair went up, an assumption corrected early in a learning review, and holds the measurement conversation described in *Where the Flow Stops*, against the history fixed on day ten: what changed, what records would show it, what moves the number for reasons nobody controls, what second measure would expose gaming, who calculates and how it is checked. Finance leads it, and the output is one page signed by the people in the room.

The decision that changes the result is whether the first reading is allowed to embarrass anyone. It will usually show that the leadership team has approved one strategy and described another. If the team treats that as a finding about its rules, it will read the next quarter with interest. If it treats it as a finding about who approved what, the lists will be tidied before the next reading.

What goes wrong. The calendars are left out. Every divergence is answered *stop*, or every one *adopt*. The unwritten rule is written down and left unchanged because changing it means refusing sensible requests. The measurement conversation is held on a gain whose history was not fixed in time; if so, agree the measure from today, and do not count the earlier improvement.

What this step does not settle is how a measured gain is then shared with the people who found it: the proportions, the timing, the thresholds and the terms. Those depend on the organisation, its people and its history. They are decided with the people who will be covered, and for an organisation working with Luminous Prosperity they are the substance of the engagement, not of a book. What the first ninety days can do is make that later decision honest, by producing a gain that both sides already believe, measured against a history neither side could move.

Day ninety: reading what you have

The working group meets on day ninety with the sponsor and seven things on the table: the first-guess page from day ten, the floor reading, the settled priority and its back-briefs, the returned balls and their answers, the pairs with their weeks of numbers, the three review pages, and the realised-strategy page with its answers and the signed measurement page.

It asks a short list of questions, in order. Where was the first guess right, and where wrong? Does the floor now describe the settled priority as a choice, in its own words? Which targets changed because a ball came back, and what did the level above have to decide? Where did a pair separate, and what did the team learn? Which assumptions changed in a review before they would otherwise have been discovered? Where do the written and realised strategies still disagree, and which answer did each divergence receive? Is there a gain

both sides believe? And then the question that matters most for the next ninety days: who has started sending bad news early, and who has stopped?

What not to conclude is as important. Ninety days are too short to say that the strategy has been deployed; they are long enough to say whether the practices are being kept for one priority. A single month in which a pair moved together is not yet a gain. And an absence of findings is itself a finding. It usually means that one of the earlier practices is not yet real, and the first place to look is the answer given to the first back-brief that disagreed.

The next ninety days take a second priority through the same order, and the quarterly strategy review adopts the form the monthly review has tested. The practices do not need to be expanded faster than that, and they are more likely to last if they are not.

What to do on Monday

If the whole of this chapter is too much to begin with, three things can be done on the first Monday, and none of them needs anyone's permission.

Ask three people who work within two levels of the floor what the strategy asks them to do differently on Monday, and write down exactly what they say.

Take one priority on the card and one ordinary case from last month, and write, before you ask anyone else, what you think the priority means for that case. Then ask one colleague on the leadership team to do the same, without seeing yours.

And at your next review, ask the owner of one green item the four questions, and keep the last five minutes for the only question that matters: what, if anything, should change in the strategy because of what you have just heard.

Santa Anna's canon describes the kind of capacity these three things build, writing of the moment in a meeting when a room could either deflect or tell the truth: "It requires no formal authority, no budget, no strategic plan." The same is true of the first Monday. It requires a notebook, three conversations, and the willingness to hear an answer that does not match the card.

Monday, day ninety-one

It is twenty past seven on a Monday in May, and the chief operating officer is at her desk before anyone else on the floor. The card is still on her desk, with two sentences in her own hand written under the first priority. In the monthly review pack that went out on Friday, each key result for that priority has a second number beside it, the same size. A page of changed assumptions from the last review lies under her coffee

cup. In the finance director's drawer there is a signed page with a measure, a history and two adjustments, and beside it a page headed *adopt, stop or test*, filled in.

None of these has a name. Nobody has launched anything. She looks at the first meeting of the week, notes what she is least sure the strategy means for it, and closes the notebook.

CHAPTER SEVEN

The Field Manual

Fifteen working pages, the readings, the run sheets, the pairs and the three answers, ready to carry into a room.



The Field Manual

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SEVEN

How to use these pages

The chapters of this book describe five practices as people doing them.[46] This chapter holds the working pages: the readings, the run sheets, the question sets and the one-page templates that the earlier chapters point to and do not print. Each page is meant to be photocopied, filled in by hand and kept in a folder that the convenor of Chapter 6 looks after. Nothing here needs software.

The pages follow the order of the first ninety days, not the order of the chapters. Each opens with a line saying where it comes from and when it is used. Where a page gives a script, the words are a starting point;

say them in your own voice. Where it gives a time, the time is a planning guide from practice, not a finding.

Two rules govern the whole manual. Write down what people said, not what you wish they had said. And answer everything you collect: a back-brief unanswered, a ball not caught, a pair never read, a page of changed assumptions that the leadership team never sees, teaches the opposite of what the practice is for.

What these pages do not contain is the design of any arrangement for sharing a gain: proportions, thresholds, timing or terms. Those belong to each organisation and the people it covers, and they are decided separately.

1 · The floor reading

From Chapter 1. Used in days one to ten, and again at day ninety.

Ask ten people, privately and without notice, at least half of them within two levels of the work: *What does the strategy ask you to do differently on Monday?* Write the answer word for word. Then tick one column.

PERSON (ROLE ONLY)	THEIR WORDS	CARD REPEATED	A MANAGER'S INSTRUCTION	A REAL CHOICE, WITH AN EXAMPLE

Reading it. Count the last column. It is a reading taken by hand, not a statistic, and it should not be reported as one. Keep the words; at day ninety, ask the same ten people the same question and lay the two pages side by side.

2 · The truce test

From Chapter 1. One meeting of the leadership team, about ninety minutes, in days one to ten.

For each priority on the card, the convenor brings one ordinary case from last month, from the floor reading where possible. Each executive writes privately, in two sentences, what the priority means for that case. Then the answers are read aloud together.

PRIORITY	THE CASE	ANSWERS THAT AGREE	ANSWERS THAT DIFFER, AND HOW

Reading it. Where the answers differ, the priority is still a truce. Do not resolve it in this meeting; write the difference down. The priority with the widest difference is the one to settle first (Chapter 6, days eleven to thirty).

Settling it. In a later meeting, with the case in front of the team, write two or three sentences: which good thing governs when the two conflict, and how much risk the firm will carry to honour it. Test the sentences on a second case before they leave the room.

3 · The back-brief page

From Chapter 1. Written by each manager two levels below the leadership team, for one settled priority, in days eleven to thirty. The sponsor writes hers first, upward.

Priority: Team: Written by: Date:

1. What this priority asks of my team, in a sentence a new starter would understand:
2. The one or two decisions it changes, as if-then sentences, each with a real case from last month:

- *If* _____ *then* _____
- *If* _____ *then* _____

1. What my team will stop doing, or do less of, because of it:
2. What I need from above to do it (an authority, a piece of information, a decision only someone else can make):

About two hundred words. No score, no milestones, no request for budget.

4 · Reading the back-briefs: run sheet

From Chapter 1. For the person who receives them.

Before replying to anyone. Read every page. Sort them by the choice each makes for the same kind of case.

THE CHOICE THE PAGE MAKES	PAGES (COUNT)	TEAMS

Within the week. Answer every page in writing, even if the answer is only that the question has gone up and when it will come back. Take the disagreements upward unedited, with their cases attached, and ask for the decision none of the writers can make.

Mistakes to watch for. Grading or correcting the wording. Answering the easy pages and leaving the awkward. Asking for back-briefs without writing your own. Treating the first page that disagrees as a misunderstanding.

Signs it is working. The second round of pages is less alike than the first round of the floor reading was, and more specific. The line on what will stop is filled in on most pages.



5 · Catchball: the three questions and the return

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SEVEN

5 · Catchball: the three questions and the return

From Chapter 2. Used in days twenty-one to forty-five, for the targets that serve the settled priority.

What goes down with each target. The target; the part of the strategy it serves, in a sentence; what the level above believes about how it could be reached.

The three questions, answered on one page by the level below, within a week.

1. What would it take to reach this target with what we have now?

2. What could we reach if one thing changed, and what is that thing?
3. What would reaching it cost elsewhere, in our own words?

The return. Keep the three apart.

What we commit to
What we could reach with a named decision from above (name it)
What we believe cannot be reached without a cost to something else the strategy values (name the cost)
What we compared with our peers before returning this

Before returning. The managers at the same level meet for an hour, without the person above them, to compare what they have heard.

6 · Catching the ball: questions for the level above

From Chapter 2. For the director who receives the returns.

Read every return before answering any. Then, for each target, ask:

- Do the returns say something about the strategy, not only about the number?
- Can the target and the strategy both be true? If not, which gives way?
- Do you genuinely have the capacity to fulfil this, and what would need to be true for it to be realistic?
- What are we willing to let go of to make space for it?
- What decision is ours to make, and by when?

Answer in the same form the return used: in writing, specific, within the week. Two rounds are enough for most targets and a third for one or two. More than three usually means the level above has not yet decided what it will let go of.

Mistakes to watch for. Restating the original target. Treating the returns as a negotiation over a shortfall. Running catchball only on targets already promised to the board.

7 · The pair page

From Chapter 3. Written for each key result on the day the target is agreed, in days thirty-one to sixty.

Purpose, in a sentence a new starter would understand
Key result
Kind (committed · aspirational · learning goal)

Paired measure: the number that would rise if the key result were met the wrong way
The record it comes from, and who keeps it
We will not …
Signed by, and date

Finding the second number. Ask: *if a sensible, tired person wanted to hit this number without doing what it is for, what would they do?* Then ask what record would show it. Common pairs, from Chapter 3:

FIRST NUMBER	A SECOND NUMBER THAT OFTEN EXPOSES IT
Handling or processing time	Repeat contacts about the same fault
Sales	Returns, discounts or cancellations the next quarter
Cost	Overtime, agency spending, faults found later and elsewhere
A delivery date	Rework after the date

Showing the pair. Both numbers, the same size, in the same place, every time either is shown.

8 · Choosing the kind of goal

From Chapter 3, after Seijts and Latham (2005) and the committed and aspirational results of Doerr (2018).

IF	THEN SET	AND REVIEW IT BY ASKING
The team knows how to do the work and the effort is what is missing	A committed result	Will we meet it, and if not, when did we first know?
The target is set high on purpose to stretch what seems possible	An aspirational result, said to be one	How far did we get, and what did we learn on the way?
The team does not yet know how to do the work well	A learning goal: find and write down a number of ways that work, by a date	Which ways did we find, and which will we keep?

The weekly review of the pair (ten minutes).

WEEK	FIRST NUMBER	SECOND NUMBER	DID THEY MOVE TOGETHER?	IF NOT, WHAT DO WE THINK IT MEANS?

Reading it. First up and second up together: probably a gain. First up, second worse: not yet a gain; find out why. First worse, second better: the team may have found something that matters more than the target. Watch especially the last weeks before a target is assessed, and the band just short of the line.

9 · The learning review: run sheet

From Chapter 4. Begun in one monthly operating review from day forty-one; the quarterly review adopts it later.

Three days before (pre-read). The colours and figures go out in writing. Colours are set by each owner together with the person nearest the work. Every owner is told, in writing, that an early amber will be thanked and never questioned on the day it appears. The meeting does not repeat the pre-read.

Choosing three (five minutes, by the chair, beforehand). One amber. One green that has moved its milestones or its scope. One green that seems to be going well. Rotate them, so that being chosen is ordinary.

For each of the three (ten to fifteen minutes). The owner answers, in order:

1. What did the strategy expect by now? (From the original plan, not the current one.)
2. What has actually happened? (Facts and dates.)
3. Why is there a difference? (Name an assumption, not a person.)
4. What would you change? (The owner answers first, before the room.)

The last fifteen minutes. What in the strategy itself changes because of what we heard? Write it on page 10.

Mistakes to watch for. Walking through the pre-read. Choosing only the troubled items. Letting question 3 become *who*. An empty page 10 for two reviews running.



10 · The colour rule and the page of changed assumptions

HAUTE LUMIÈRE · THE STRATEGY THAT REACHES THE FLOOR · CHAPTER SEVEN

10 · The colour rule and the page of changed assumptions

From Chapter 4.

The colour rule. Set by the owner and the person nearest the work together. An early amber: thanked, and helped. A late amber: one question, asked without heat, *when did you first know?* Never the same reception for both.

The page that leaves the room.

ASSUMPTION IN THE STRATEGY	WHAT WE NOW BELIEVE	WHAT CHANGES BECAUSE OF IT	OWNER	BY WHEN

The page goes to whoever owns the strategy, with the colours. If it is empty, write *nothing changed* and the date; two empty pages in a row mean the review is not yet real.

11 · The three lists and the realised-strategy page

From Chapter 5. Assembled for the previous quarter in days sixty-one to ninety, then every quarter.

The three lists. From records the organisation already keeps, sorted by two people who did not make the decisions.

PRIORITY ON THE SLIDE	MONEY APPROVED	NEW ROLES FILLED	LEADERSHIP HOURS (FROM CALENDARS)

None of these

The meeting (about two hours). Spend it only on the places where the lists and the slide disagree. For each, one answer on the record:

DIVERGENCE	ADOPT (THE PATTERN IS WISER; THE STRATEGY CHANGES)	STOP (THE PATTERN IS DRIFT; THE RULE CHANGES)	TEST (NOBODY KNOWS; DESIGN AN EXPERIMENT)	OWNER AND DATE

A test, on one line. We believe ; we will know by (date) because (measure); we will stop if .

Reading it. Over a year, a healthy team uses all three answers, and can point to one place where its written strategy changed because its realised one was wiser.

12 · Finding the unwritten rule

From Chapter 5. About an hour, before the meeting on page 11.

Take the last twenty approvals and the last twenty refusals or deferrals. Ask the people who decided them for the reason, as they remember it. Write the reasons in their words and count them.

THE REASON, IN THE DECIDER'S WORDS	TIMES GIVEN	SERVES WHICH PRIORITY?

Write the two or three most frequent reasons at the top of page 11, beside the strategy, in the same size of type. Then decide: keep the rule, change it, or write a second rule beside it for the part of the business it is starving.

13 · The ninety days on one page

From Chapter 6.

DAYS	PRACTICE	WHO	TIME IT TAKES	WHAT IT PRODUCES	EARLY SIGN IT IS GOING WRONG
1–10	Read: the floor reading, the truce test, the list of targets and reviews, and a dated copy of the history (pages 1 and 2)	Sponsor, convenor, finance	Sponsor about 5 hours a week; one 90-minute leadership meeting	A first-guess page	People asking what the new initiative is called
11–30	Settle one priority; the back-brief in one division (pages 2 to 4)	Leadership team; managers two levels down	One meeting to settle; a week to write; a week to answer	Settled sentences; pages with what will stop	The back-briefs are more alike than the floor reading
21–45	Catchball on that priority's targets (pages 5 and 6)	Director and the level below	A week a round; two rounds	Returned targets; decisions made above; costs named	No return names a cost elsewhere
31–60	The pair for each key result (pages 7 and 8)	Each team holding a key result	An hour to write; ten minutes a week	Pairs, with weeks of both numbers	The second number not looked at by week three
41–75	The learning review in one monthly meeting (pages 9 and 10)	Chair; owners; people nearest the work	About an hour, three times	Three pages of changed assumptions	Page 10 empty twice
61–90	The realised-strategy reading; the unwritten rule; one measurement conversation (pages 11 and 12; *Where the Flow Stops*, Chapter 7)	Working group; leadership team; finance	A day to assemble; a two-hour meeting; two 90-minute meetings to measure	Adopt, stop or test on the record; a signed measurement page	The calendars left out

14 · The day-ninety review

From Chapter 6. A meeting of the working group and the sponsor, about two hours.

Bring the first-guess page, both floor readings, the settled sentences and back-briefs, the returned balls and answers, the pairs, the three review pages, the realised-strategy page and the signed measurement page.

Ask, in order:

1. Where was the first guess right, and where was it wrong?
2. Does the floor now describe the settled priority as a choice, in its own words?
3. Which targets changed because a ball came back, and what did the level above have to decide?
4. Where did a pair separate, and what did the team learn?
5. Which assumptions changed in a review before they would otherwise have been found?
6. Where do the written and realised strategies still disagree, and what answer did each divergence receive?
7. Is there a gain both sides believe, and what does it show so far?
8. Who has started sending bad news early, and who has stopped?

What not to conclude. That the strategy has been deployed (ninety days show whether the practices are kept for one priority). That one month of a pair moving together is a gain. That an absence of findings means an absence of problems; look first at the answer given to the first back-brief that disagreed.

15 · The first Monday

From Chapter 6. For anyone who wants to begin with less.

- Ask three people within two levels of the floor what the strategy asks them to do differently on Monday. Write down exactly what they say.
- Take one priority and one ordinary case from last month. Write what you think the priority means for it, then ask one colleague on the leadership team to do the same without seeing yours.
- At your next review, ask the owner of one green item the four questions, and keep the last five minutes for what should change in the strategy.

The plan on your own floor

A strategy reaches the floor when it tells one person which of two good things to choose on an ordinary morning.

Luminous Prosperity Inc. takes this work into organisations as a gainshare. The flow is opened from the inside, with the people who already hold what the firm knows. The gain is counted from records the firm already keeps, against a measure both sides sign before anyone knows the answer, and the house is paid as a share of what is found: between thirty and ten per cent of it.

The book gives the card, the pairs and the three answers. How your own plan reads on a Monday morning, and what closing the distance would return, is read from your people and your records. That begins with a conversation.

REQUEST A CONVERSATION

References

1. The author's canon in this chapter, and the works it draws on: Santa Anna, A. *Conscious Governance Models*, "The Cost of Unconscious Governance" and "Principle 2: Decisions Are Made at the Closest Point to the Relevant Intelligence". The Luminous Developmental Canon; Santa Anna, A. *Integrated Value Stream Mapping*, "Principle 2: Positional Diversity". The Luminous Developmental Canon; Santa Anna, A. *The Luminous 100*, "Balanced Scorecard (Kaplan & Norton)". Luminous Prosperity Inc; Santa Anna, A. *Organizational Soul Work*, "The Room Before the Room" and "Case Study: The Emergence at Meridian" (a composite in the canon). The Luminous Developmental Canon; Santa Anna, A. *The Keystone Effect*, "Step 2: Map the Information Flow". The Luminous Developmental Canon; Santa Anna, A. *Psychological Safety Architecture*, "Signs of Dorsal Vagal Shutdown (Collapse)". The Luminous Developmental Canon.
2. Sull, D., Homkes, R., & Sull, C. (2015). Why strategy execution unravels, and what to do about it. *Harvard Business Review*, 93(3), 57–66. (No Crossref DOI. The original text was not reachable for this issue; the figures are as reported in contemporary summaries of the article, which agree with each other. Survey evidence in a practitioner magazine; graded as such in the prose.)
3. Hambrick, D. C., & Cannella, A. A. (1989). Strategy implementation as substance and selling. *Academy of Management Executive*, 3(4), 278–285. https://doi.org/10.5465/ame.1989.4277401 (Argument summarised; full text not reopened.)
4. Kaplan, R. S., & Norton, D. P. (1996). Linking the balanced scorecard to strategy. *California Management Review*, 39(1), 53–79. https://doi.org/10.2307/41165876 (Argument summarised; full text not reopened.)
5. Kaplan, R. S., & Norton, D. P. (2004). *Strategy maps: Converting intangible assets into tangible outcomes*. Harvard Business School Press. (Book; verified in Open Library.)
6. Bungay, S. (2011). *The art of action: How leaders close the gaps between plans, actions and results*. Nicholas Brealey. (Book; verified in Open Library. The three gaps and the back-brief are summarised from memory of the book; its text was not reopened.)
7. Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 493–503. https://doi.org/10.1037/0003-066X.54.7.493
8. Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119. https://doi.org/10.1016/S0065-2601(06)38002-1 (Ninety-four tests and the effect size, d = .65, from the abstract.)
9. The author's canon in this chapter, and the works it draws on: Santa Anna, A. *Integrated Value Stream Mapping*, on positional bias in who maps a value stream. The Luminous Developmental Canon; Santa Anna, A. *Organizational Soul Work*, on the differential cost of silence. The Luminous Developmental Canon; Santa Anna, A. *Conscious Governance Models*, on information distortion in concentrated authority, and "Movement 2: Commitment — Making Promises We Mean". The Luminous Developmental Canon; Santa Anna, A. *The Luminous 100*, "OKRs (Objectives and Key Results)". Luminous Prosperity Inc.
10. Guth, W. D., & MacMillan, I. C. (1986). Strategy implementation versus middle management self-interest. *Strategic Management Journal*, 7(4), 313–327. https://doi.org/10.1002/smj.4250070403 ("Counter effort" and the redirect, delay, reduce and sabotage findings are from the abstract.)
11. Floyd, S. W., & Wooldridge, B. (1992). Middle management involvement in strategy and its association with strategic type: A research note. *Strategic Management Journal*, 13(S1), 153–167. https://doi.org/10.1002/smj.4250131012 (The four roles, upward and downward, summarised; full text not reopened.)
12. Floyd, S. W., & Wooldridge, B. (1997). Middle management's strategic influence and organizational performance. *Journal of Management Studies*, 34(3), 465–485. https://doi.org/10.1111/1467-6486.00059 (The sample of 259 and both findings are from the abstract.)
13. Wooldridge, B., Schmid, T., & Floyd, S. W. (2008). The middle management perspective on strategy process: Contributions, synthesis, and future research. *Journal of Management*, 34(6), 1190–1221. https://doi.org/10.1177/0149206308324326 (Review; summarised.)
14. Westley, F. R. (1990). Middle managers and strategy: Microdynamics of inclusion. *Strategic Management Journal*, 11(5), 337–351. https://doi.org/10.1002/smj.4250110502

15. Huy, Q. N. (2002). Emotional balancing of organizational continuity and radical change: The contribution of middle managers. *Administrative Science Quarterly*, 47(1), 31–69. https://doi.org/10.2307/3094890

16. Balogun, J., & Johnson, G. (2004). Organizational restructuring and middle manager sensemaking. *Academy of Management Journal*, 47(4), 523–549. https://doi.org/10.5465/20159600 (Findings summarised from memory of the published study; no abstract was retrievable and the text was not reopened.)

17. Akao, Y. (Ed.). (1991). *Hoshin kanri: Policy deployment for successful TQM*. Productivity Press. (Book; verified in Open Library. The 1960s origin is the standard account in the practitioner literature, not checked against a primary Japanese source.)

18. Witcher, B., & Butterworth, R. (2001). Hoshin kanri: Policy management in Japanese-owned UK subsidiaries. *Journal of Management Studies*, 38(5), 651–674. https://doi.org/10.1111/1467-6486.00253 (The two-year project, the three subsidiaries, QCDE, and the difficulties named are from the abstract.)

19. Tennant, C., & Roberts, P. (2001). Hoshin kanri: Implementing the catchball process. *Long Range Planning*, 34(3), 287–308. https://doi.org/10.1016/S0024-6301(01)00039-5 (Cited for the subject of the study; its findings were not reopened.)

20. The author's canon in this chapter, and the works it draws on: Santa Anna, A. *Values-Based Performance Systems*, "The Evolution of the Measurement Paradigm" and "What Gets Measured Gets Managed — And What Gets Lost", "The Somatic Cost of Misaligned Measurement" and "Case Study: The Organization That Optimized Itself Into Hollowness" (a composite in the canon). The Luminous Developmental Canon; Santa Anna, A. *Conscious Governance Models*, on fear-based accountability ("Sacred Accountability"). The Luminous Developmental Canon.

21. Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, 57(9), 705–717. https://doi.org/10.1037/0003-066X.57.9.705 (The core findings, mechanisms and moderators are as the abstract lists them; the detail of the complex-task moderator is summarised from the theory's standard statement, not reopened.)

22. Seijts, G. H., & Latham, G. P. (2005). Learning versus performance goals: When should each be used? *Academy of Management Executive*, 19(1), 124–131. https://doi.org/10.5465/ame.2005.15841964

23. Ordóñez, L. D., Schweitzer, M. E., Galinsky, A. D., & Bazerman, M. H. (2009). Goals gone wild: The systematic side effects of overprescribing goal setting. *Academy of Management Perspectives*, 23(1), 6–16. https://doi.org/10.5465/amp.2009.37007999

24. Locke, E. A., & Latham, G. P. (2009). Has goal setting gone wild, or have its attackers abandoned good scholarship? *Academy of Management Perspectives*, 23(1), 17–23. https://doi.org/10.5465/amp.2009.37008000

25. Schweitzer, M. E., Ordóñez, L., & Douma, B. (2004). Goal setting as a motivator of unethical behavior. *Academy of Management Journal*, 47(3), 422–432. https://doi.org/10.5465/20159591 (That the participants were students is from memory of the study, not from the abstract.)

26. Bevan, G., & Hood, C. (2006). What's measured is what matters: Targets and gaming in the English public health care system. *Public Administration*, 84(3), 517–538. https://doi.org/10.1111/j.1467-9299.2006.00600.x (Targets and terror, the Soviet parallel and synecdoche are from the abstract; that the regime also produced real improvements is the authors' wider argument, summarised.)

27. Kerr, S. (1975). On the folly of rewarding A, while hoping for B. *Academy of Management Journal*, 18(4), 769–783. https://doi.org/10.5465/255378

28. Grove, A. S. (1983). *High output management*. Random House. (Book; verified in Open Library. Paired indicators summarised from memory; the text was not reopened.)

29. Doerr, J. (2018). *Measure what matters*. Portfolio. (Book; verified in Open Library. Committed and aspirational key results summarised from memory; the text was not reopened.)

30. The author's canon in this chapter, and the works it draws on: Santa Anna, A. *Organizational Soul Work*, "Type 1: Operational Truths — 'This Isn't Working'". The Luminous Developmental Canon; Santa Anna, A. *Luminous Living Systems*, "Domain 4: Feedback Loop Vitality — The Integrity of Information Flow". The Luminous Developmental Canon; Santa Anna, A. *Conscious Governance Models*, on information distortion in concentrated authority; "Principle 2: Experiment Rather Than Execute"; "The Word We Fear" and "Movement 2: Commitment — Making Promises We Mean" (Chapter 10, Sacred Accountability); and "The Governance Retrospective". The Luminous Developmental Canon.

31. Sull, D., Homkes, R., & Sull, C. (2015). Why strategy execution unravels, and what to do about it. *Harvard Business Review*, 93(3), 57–66. (No Crossref DOI. The original text was not reachable; the myth of sticking to the plan and the findings on adapting are as reported in contemporary summaries of the article. Survey evidence; graded in the prose.)
32. The four questions and the army's training centres are the standard description of the after-action review in the practitioner literature, summarised here; no single primary army document was opened for this issue.
33. Darling, M. J., & Parry, C. S. (2001). After-action reviews: Linking reflection and planning in a learning practice. *Reflections: The SoL Journal*, 3(2), 64–72. https://doi.org/10.1162/152417301526952 52 (Practitioner article; cited for the description of the practice; text not reopened.)
34. Tannenbaum, S. I., & Cerasoli, C. P. (2013). Do team and individual debriefs enhance performance? A meta-analysis. *Human Factors*, 55(1), 231–245. https://doi.org/10.1177/0018720812448394 (Forty-six samples, N = 2,136, about 25 per cent, d = .67, and the moderators are from the abstract.)
35. Ellis, S., & Davidi, I. (2005). After-event reviews: Drawing lessons from successful and failed experience. *Journal of Applied Psychology*, 90(5), 857–871. https://doi.org/10.1037/0021-9010.90.5.857
36. Shewhart, W. A. (1931). *Economic control of quality of manufactured product*. Van Nostrand; and Deming, W. E. (1986). *Out of the crisis*. MIT Center for Advanced Engineering Study (Open Library dates its first issue to 1982). (Books; the attribution of the plan-do-check-act cycle to the two is the standard account, summarised from memory, not reopened.)
37. Witcher, B., & Butterworth, R. (2001). Hoshin kanri: Policy management in Japanese-owned UK subsidiaries. *Journal of Management Studies*, 38(5), 651–674. https://doi.org/10.1111/1467-6486.00253 (Review as the driver, and the difficulty of administering periodic review, are from the abstract.)
38. The author's canon in this chapter, and the works it draws on: Santa Anna, A. *Conscious Governance Models*, Chapter 9, "Emergent Strategy — Governing at the Edge of What We Cannot Yet Know": "The Illusion of the Master Plan", "The Origins of Emergent Strategy", "Principle 2: Experiment Rather Than Execute" and "Principle 4: Hold Direction Lightly". The Luminous Developmental Canon. (The cooperative's retreat is the canon's own illustrative account.); Santa Anna, A. *Luminous Living Systems: The Organization as Conscious Organism*, on probe, sense and respond and on nonlinearity; and *The Luminous 100*, "Cynefin Framework (Dave Snowden)". The Luminous Developmental Canon; Luminous Prosperity Inc; Santa Anna, A. *Adaptive Resource Orchestration*, "What Living Systems Theory Illuminates". The Luminous Developmental Canon.
39. Mintzberg, H. (1978). Patterns in strategy formation. *Management Science*, 24(9), 934–948. https://doi.org/10.1287/mnsc.24.9.934 (The definition, the two studies and the conclusion are from the abstract.)
40. Mintzberg, H., & Waters, J. A. (1985). Of strategies, deliberate and emergent. *Strategic Management Journal*, 6(3), 257–272. https://doi.org/10.1002/smj.4250060306 (The continuum and the named types are from the abstract; the descriptions of umbrella and imposed strategies are summarised from the paper's standard reading, not reopened.)
41. Burgelman, R. A. (1991). Intraorganizational ecology of strategy making and organizational adaptation: Theory and field research. *Organization Science*, 2(3), 239–262. https://doi.org/10.1287/orsc.2.3.239 (Induced and autonomous processes and the proposition on reorientation are from the abstract.)
42. Burgelman, R. A. (1994). Fading memories: A process theory of strategic business exit in dynamic environments. *Administrative Science Quarterly*, 39(1), 24–56. https://doi.org/10.2307/2393493 (No abstract was retrievable and the full text was not reopened; the allocation rule and the sequence of the exit are as described in secondary accounts of Burgelman's Intel research, which differ on some dates.)
43. Grove, A. S. (1996). *Only the paranoid survive*. Currency Doubleday. (Book; verified in Open Library; not reopened.)
44. Bower, J. L. (1970). *Managing the resource allocation process: A study of corporate planning and investment*. Harvard University, Graduate School of Business Administration. (Book; verified in Open Library; summarised from its standard reading, not reopened.)
45. Sources for this chapter: The practices and their sources are those of Chapters 1 to 5: Sull, Homkes and Sull, Hambrick and Cannella, Kaplan and Norton, Bungay and Gollwitzer on communication and translation; Guth and MacMillan, Floyd and Wooldridge, Westley, Huy, Balogun and Johnson, Akao, Witcher and Butterworth, and Tennant and Roberts on the middle and catchball; Locke and Latham, Seijts and Latham, Ordóñez and colleagues, Schweitzer and colleagues, Bevan and Hood, Kerr,

Grove and Doerr on goals and their measures; Tannenbaum and Cerasoli, Ellis and Davidi, and Witcher and Butterworth on review; Mintzberg, Mintzberg and Waters, Burgelman and Bower on realised strategy. Full references are in those chapters' notes; the measurement conversation is set out in Santa Anna, A. (2026), *Where the Flow Stops*, Chapter 5 and Chapter 7; Santa Anna, A. *Conscious Governance Models*, "Why More Process Is Not the Answer", "Principle 1: Sense Before You Strategize", "Principle 2: Experiment Rather Than Execute" and "Principle 4: Hold Direction Lightly". The Luminous Developmental Canon; Santa Anna, A. *The Keystone Effect*, on holding the moment in a meeting. The Luminous Developmental Canon; The time estimates in this chapter are planning guides from practice, not findings of a study; adjust them to the organisation.

46. Sources for this chapter: The sources behind these pages are given in full in the notes to Chapters 1 to 5: Sull, Homkes and Sull (2015), Kaplan and Norton (1996, 2004), Bungay (2011) and Gollwitzer (1999; with Sheeran, 2006) for pages 1 to 4; Guth and MacMillan (1986), Floyd and Wooldridge (1992, 1997), Westley (1990), Akao (1991), Witcher and Butterworth (2001) and Tennant and Roberts (2001) for pages 5 and 6; Locke and Latham (2002), Seijts and Latham (2005), Schweitzer, Ordóñez and Douma (2004), Bevan and Hood (2006), Grove (1983) and Doerr (2018) for pages 7 and 8; Tannenbaum and Cerasoli (2013), Ellis and Davidi (2005) and Darling and Parry (2001) for pages 9 and 10; Mintzberg (1978), Mintzberg and Waters (1985), Burgelman (1991, 1994) and Bower (1970) for pages 11 and 12; the measurement conversation is in Santa Anna, A. (2026), *Where the Flow Stops*; Santa Anna, A. *Conscious Governance Models*, "Movement 2: Commitment — Making Promises We Mean" (the capacity questions on page 6). The Luminous Developmental Canon; The times given are planning guides from practice, not findings of a study.

The Strategy That Reaches the Floor

Ammanuel Santa Anna · Luminous Prosperity Inc. · Haute Lumière

Set in Cormorant Garamond and Manrope.

Read the Business Library at business.hautelumiere.life/library.

Luminous Library Licence 1.0 · hautelumiere.life/library-licence

REQUEST A CONVERSATION

business.hautelumiere.life/book/?kind=gainshare

The Business Library

business.hautelumiere.life/library/

The packages

hautelumiere.life/pricing/