NEUROSCIENCE & LEAN CULTURE

The Brain and the Standard

What neuroscience tells us about Lean initiatives that hold, and those that collapse

Ossama Yakhafallah

sciences.yakhafallah.com

Contents

An introduction and 11 chapters · about 1 h 31 of reading

Intro	**Introduction** What your brain does while you change things	6 min
01	**The Toyota paradox** Two memories, only one of them copies	7 min
02	**The transfer** What moves when a movement becomes automatic	8 min
03	**Why your standards fade** Four parameters, four structures	8 min
04	**Four slots, no more** Why visual management actually works	8 min
05	**The brain knows before you do** Fifty milliseconds, and nobody says a word	9 min
06	**Under pressure, the brain changes pilot** Why a stressed organisation turns rigid	9 min
07	**The work happens during the breaks** What the hippocampus replays while you do nothing	8 min
08	**The day nothing surprises anymore** Dopamine does not encode pleasure	9 min
09	**What circulates and isn't true** The chapter that should make you wary of this book	8 min
10	**What this book does not prove** Where the weaknesses are	6 min
11	**What changes on Monday** Seven questions, none about people	5 min

INTRODUCTION

Introduction

What your brain does while you change things

6 min read

You walk back through a workshop six months after an improvement project.

The floor markings are still there, slightly worn. The tracking board shows figures from seven weeks ago. The workstation runs exactly as it did before.

Nobody sabotaged anything. The team is the same, still willing, and if you ask them they'll tell you the standard was a good idea.

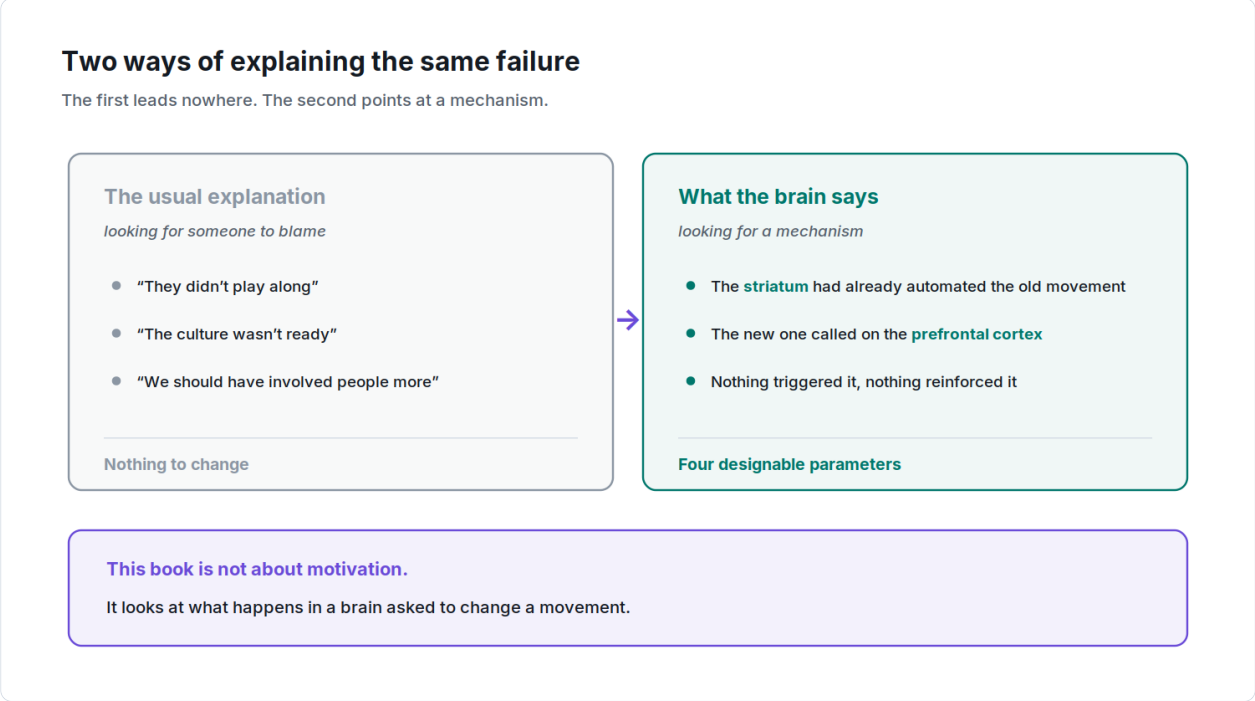
It simply faded.

The explanation you'll be given

It comes in three phrases you already know.

Lack of buy-in. The culture wasn't ready. We should have involved people more.

These explanations share one flaw: **none of them points at anything you can change**. They refer to states of mind, and a state of mind cannot be designed.



Two ways of explaining the same failure

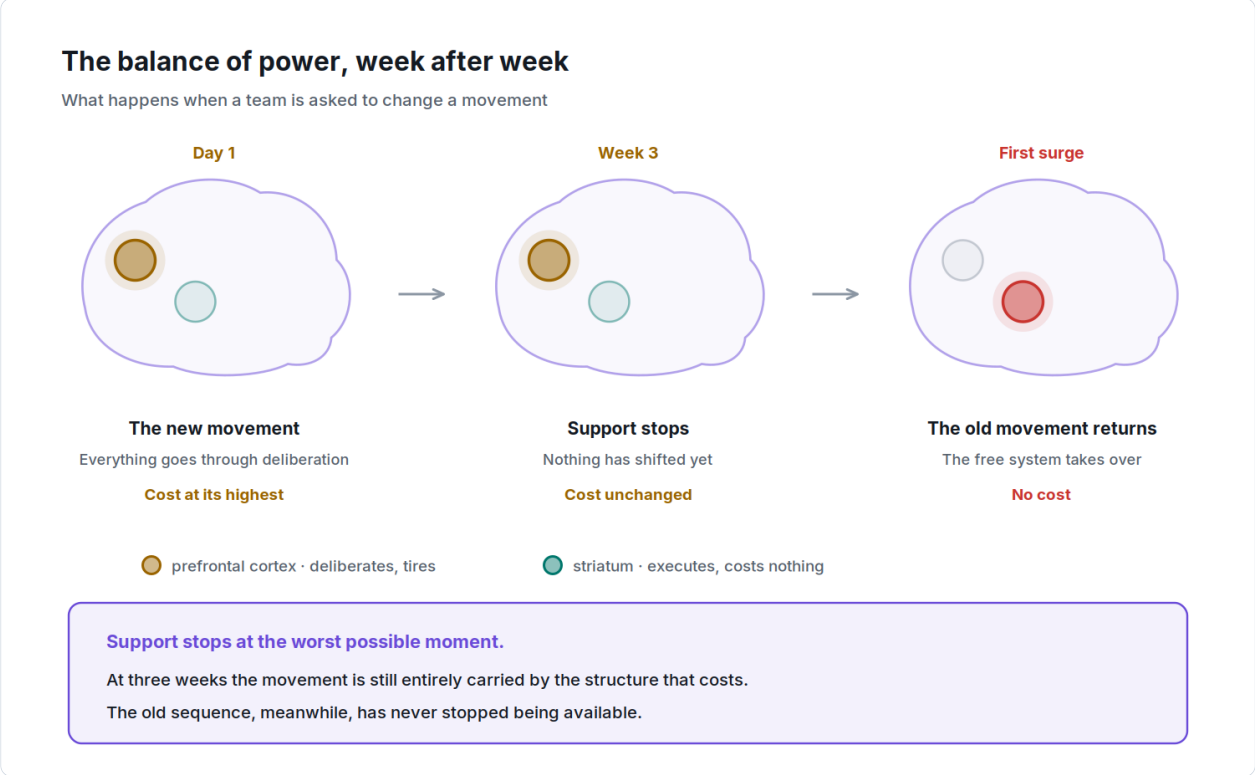
This book starts somewhere else. Not with what people want, but with what a brain does when it's asked to change a movement.

What actually happens

When a team abandons a standard after six weeks, here is the neural version of events.

The old movement was encoded in the striatum, a deep structure that runs learned sequences without consuming attention. It cost nothing.

The new standard called on the prefrontal cortex instead: the region that deliberates, compares, holds things in mind. It is powerful and slow, and above all, it tires.



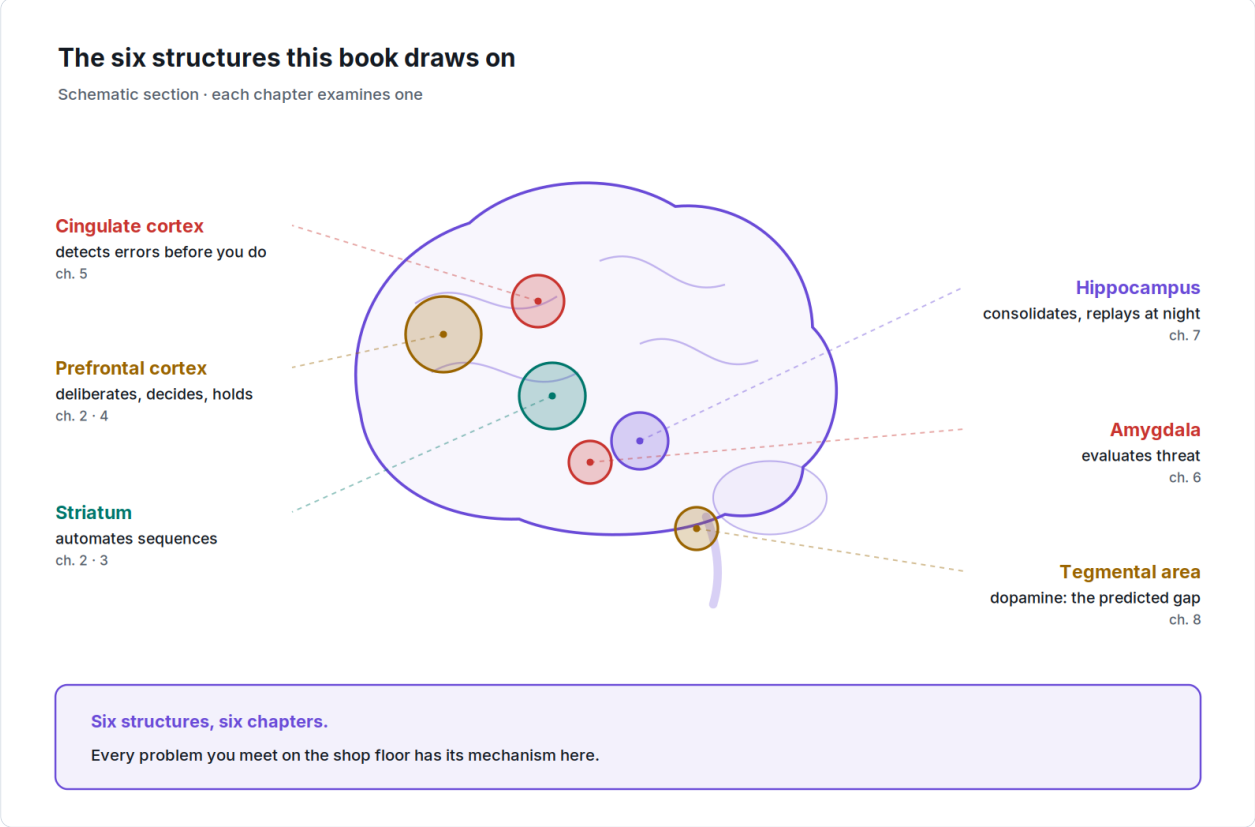
The balance of power between the two structures, week after week

On the first day of pressure, the expensive system gave way to the free one.

This isn't a character flaw. It's the expected behaviour of any human brain, including yours.

Six structures, eleven chapters

Every problem you meet on the shop floor has an identifiable mechanism. This book works through six of them.



The six brain structures this book draws on

The striatum explains why a standard fades, and why twenty-one days of support is never enough.

The prefrontal cortex explains why a seven-point standard will never be applied in full, and why visual management actually works.

The anterior cingulate cortex detects your errors **before you are aware of them**. That gap of a few hundred milliseconds is one of the most unsettling findings in this book.

The amygdala evaluates threat, and largely decides whether someone speaks up or stays silent.

The hippocampus consolidates what has been learned, mostly during sleep. Three days of training with no night in between works against it.

The ventral tegmental area releases dopamine, which does not encode pleasure, contrary to what you'll read everywhere, but the **gap between what was predicted and what arrives**. That is mechanically why the energy of an improvement project fades.

What this book is not

It is not a scientific publication. It synthesises the work of other researchers. Every claim is referenced; you can check.

It is not a "neuro-management" book. That field is crowded with shortcuts: left brain versus right brain, learning styles, profiles people claim to identify on the shop floor. A whole chapter is given to them, so you can recognise them.

It is not a method. The mechanisms described here don't produce a recipe. They produce design criteria, which is more demanding.

One caution

The work cited here comes from the laboratory: simple tasks, small samples, sometimes animals. Going from "the striatum encodes a motor sequence" to "your standard won't hold" is a considerable leap.

I make that leap, since it's the point of the book, but I flag it every time. Chapter 10 examines systematically what none of this allows us to conclude.

It should be read before applying any of what precedes it.

Let's begin with the mechanism that explains all the others: what happens when a movement stops being thought about.

CHAPTER 01

The Toyota paradox

Two memories, only one of them copies

7 min read

Since the 1980s, Toyota has opened its doors.

The company let competitors tour its plants. It published its methods. Its managers wrote books. There is, as far as I know, no protected industrial secret in its production system.

Forty years later, almost nobody has reproduced it durably.

Not for lack of means. Substantial groups, with excellent engineers, got results, then lost them.

Something doesn't transfer. And that something has a precise explanation in the brain.

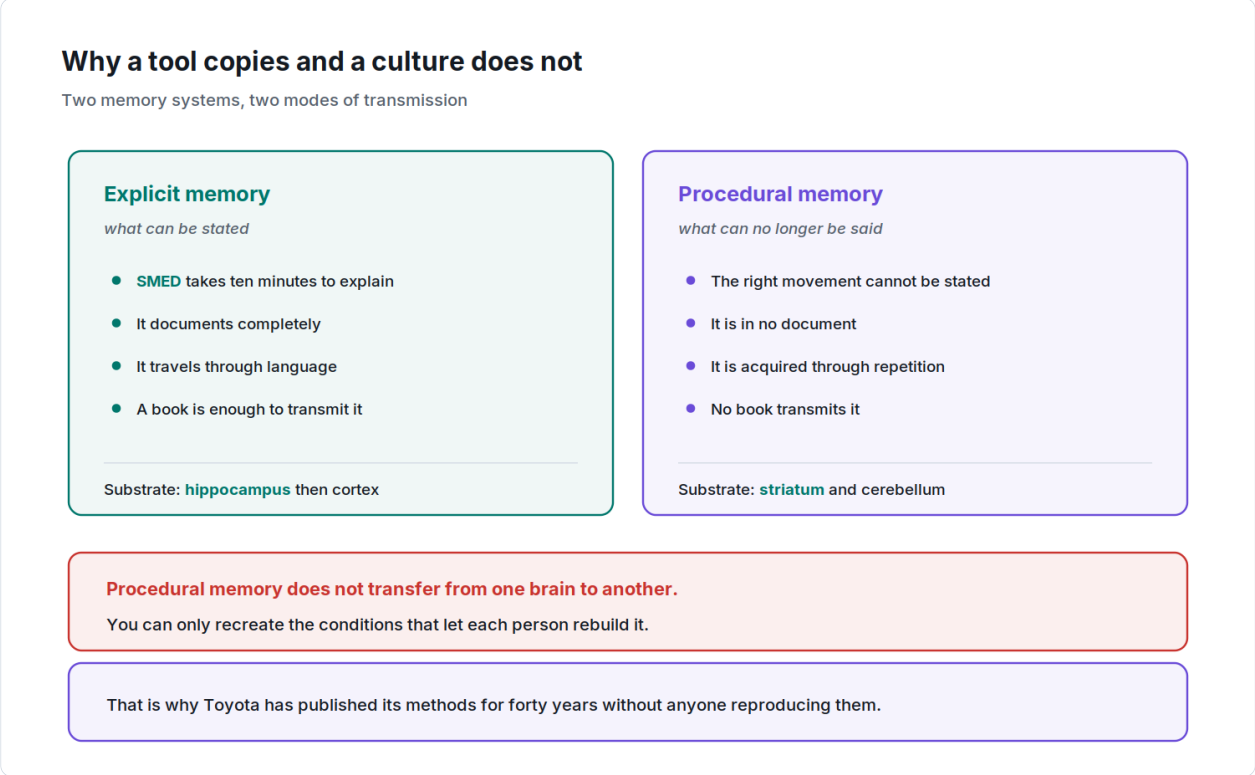
Two memories, not one

Your brain doesn't store everything in the same place. It has distinct memory systems, resting on different structures and obeying different rules.

Two concern us here.

Explicit memory holds what you can state: facts, rules, procedures you can describe. It passes through the hippocampus, which encodes it, before being gradually stored in the cortex.

Procedural memory holds what you know how to do without being able to say it. It rests on the striatum and the cerebellum. It is acquired through repetition, never through explanation.

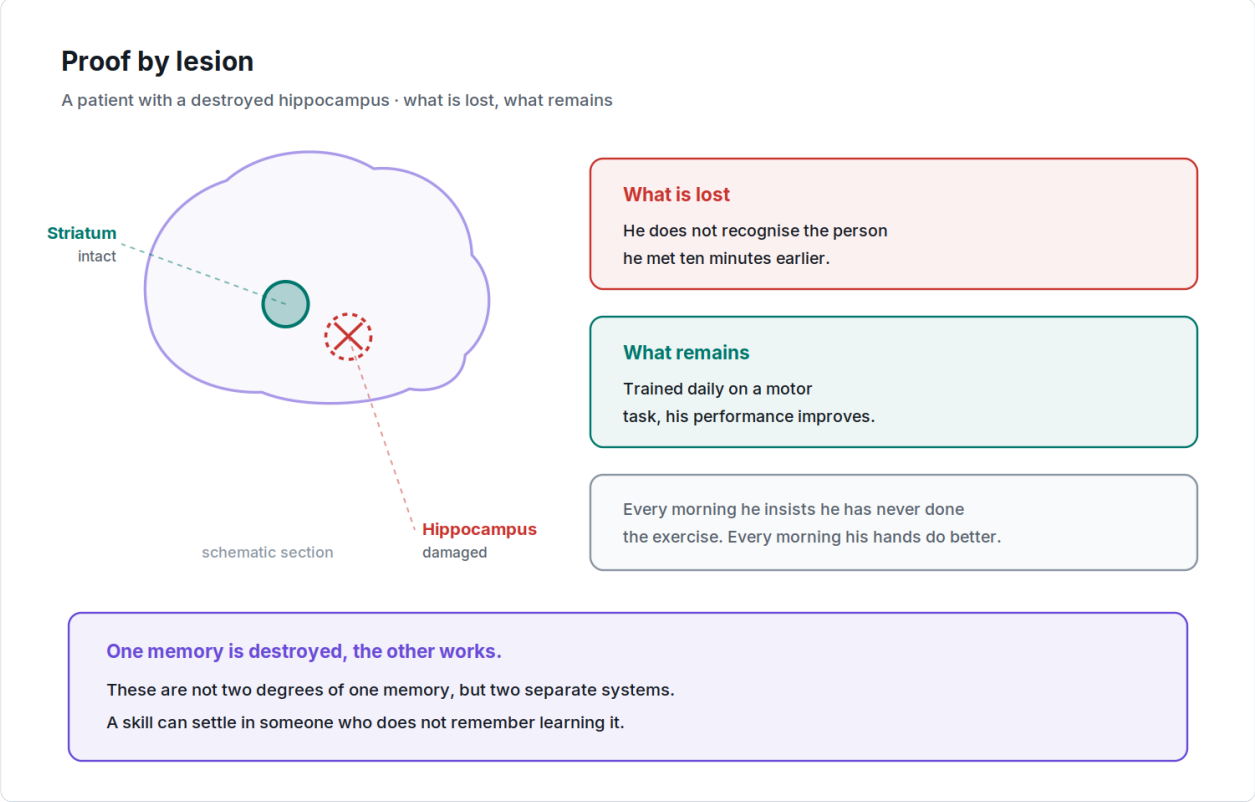


Two memory systems, two modes of transmission

The clearest demonstration of this separation comes from amnesic patients. After damage to the hippocampus, some become unable to form new memories: they don't recognise someone they met ten minutes earlier.

And yet, trained on a motor task day after day, **they improve**. Their performance progresses normally.

Every morning they insist they have never done the exercise. Every morning their hands do better than the day before.



What a patient with hippocampal damage loses and keeps

One memory is destroyed, the other works. These are not two degrees of the same thing, but two separate systems.

What this says about Lean

Let's return to the opening asymmetry with this lens.

A tool is explicit memory. SMED states itself in three sentences: separate what requires the machine stopped from what can be done while it runs, then convert. You've just learned it. You could explain it tomorrow.

A culture is collective procedural memory. It's an operator stopping the line without wondering whether they're allowed to. It's a manager whose first reflex is to ask what made the error possible. These movements cannot be stated: their authors would be unable to describe them.

And procedural memory **does not transfer from one brain to another**. It doesn't travel through language. It has to be rebuilt in each person, through repetition, under precise conditions.

That is the whole Toyota paradox. The company published its explicit memory, which is copyable. Its procedural memory stayed in the hands of its operators, where no book can go and fetch it.

The word that misleads

One term recurs in every transformation review, and it deserves questioning: **resistance**.

Look at what it actually describes.

A team has worked a certain way for years. That way is encoded in their striatum: it costs nothing, it's reliable, it works on a Friday at five in the afternoon after a difficult night.

They're asked to adopt a new standard. That standard requires, at every execution, remembering, checking, not falling back into the old movement. It calls on the prefrontal cortex, which tires.

Three weeks later, support ends. Under the first real pressure, the old sequence takes back control.

Where is the resistance?

Nobody resisted. An expensive system gave way to a free one, which is the expected behaviour of any brain, including that of the consultant recording the failure.

> *As long as you're looking for someone to blame, you aren't designing the conditions.*

What changes

Moving from "they didn't want to" to "the mechanism wasn't respected" isn't a diplomatic softening. It's what makes the problem workable.

If the problem is motivation, the levers are conviction and charisma, which don't reproduce.

If the problem is procedural memory, the levers become parameters: **the duration** of repetition required, **the trigger** that launches the sequence, **the load** the standard imposes, **the feedback** that reinforces it.

Four parameters, all designed, all independent of who implements them.

The chapters that follow take them one at a time.

What to take away

Two distinct memories. Explicit memory passes through the hippocampus and travels through language. Procedural memory rests on the striatum and is acquired only through repetition.

Amnesic patients demonstrate it. They learn movements they don't remember learning.

A tool is explicit, a culture is procedural. That's why one travels and the other doesn't.

"Resistance" names a mechanism, not an opposition. The free system prevails over the expensive one.

Sources

Squire, L. R. Work on multiple memory systems and amnesic patients, establishing the dissociation between declarative and procedural memory.

Cohen, N. J. & Squire, L. R. (1980). *Preserved learning and retention of pattern-analyzing skill in amnesia*. Science. Procedural learning preserved despite declarative amnesia.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). *A neostriatal habit learning system in humans*. Science. The role of the striatum in habit learning.

CHAPTER 02

The transfer

What moves when a movement becomes automatic

8 min read

Take someone who has driven for twenty years. Ask them to explain what they do when they change gear.

They won't be able to. They'll describe something approximate, without giving the exact sequence or the micro-adjustments they make fifty times a journey.

Put them behind the wheel of a nineteen-tonne lorry and they become able to describe everything again.

Nothing has changed in their competence. The movement has changed structure.

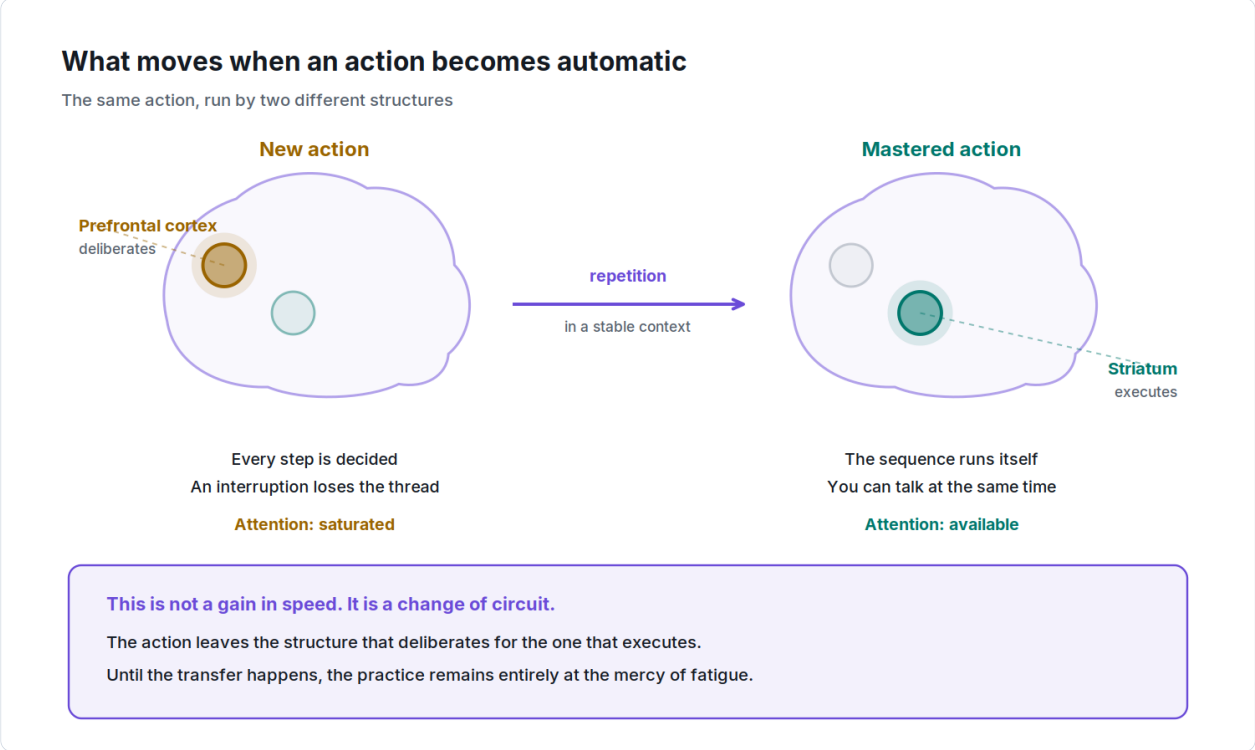
Two circuits for one movement

The brain has two routes for producing behaviour, and they coexist permanently.

The prefrontal cortex, behind your forehead, deliberates. It compares, decides, holds an instruction in mind. It is extraordinarily flexible: you can have it execute something never done before, from a simple explanation.

It has one crippling flaw. It consumes attention, it tires, it handles one thing at a time.

The striatum, a deep structure, does only one thing: replay a learned sequence as soon as a known signal appears. It is rigid, and almost free.



The same movement executed by two different structures

The shift from one to the other has a name: automation. And it cannot be decreed.

Why the brain insists on it

Your brain is about 2% of your mass and consumes about 20% of your energy at rest. Proportionally, it is the most expensive organ in the body.

A system that costly is under permanent pressure to economise. Hence a principle that explains much of how it works: **anything that can be automated eventually will be**.

Not out of laziness. Because handing a sequence to the striatum frees the scarcest resource in the organism: conscious attention.

This has a direct consequence for your work.

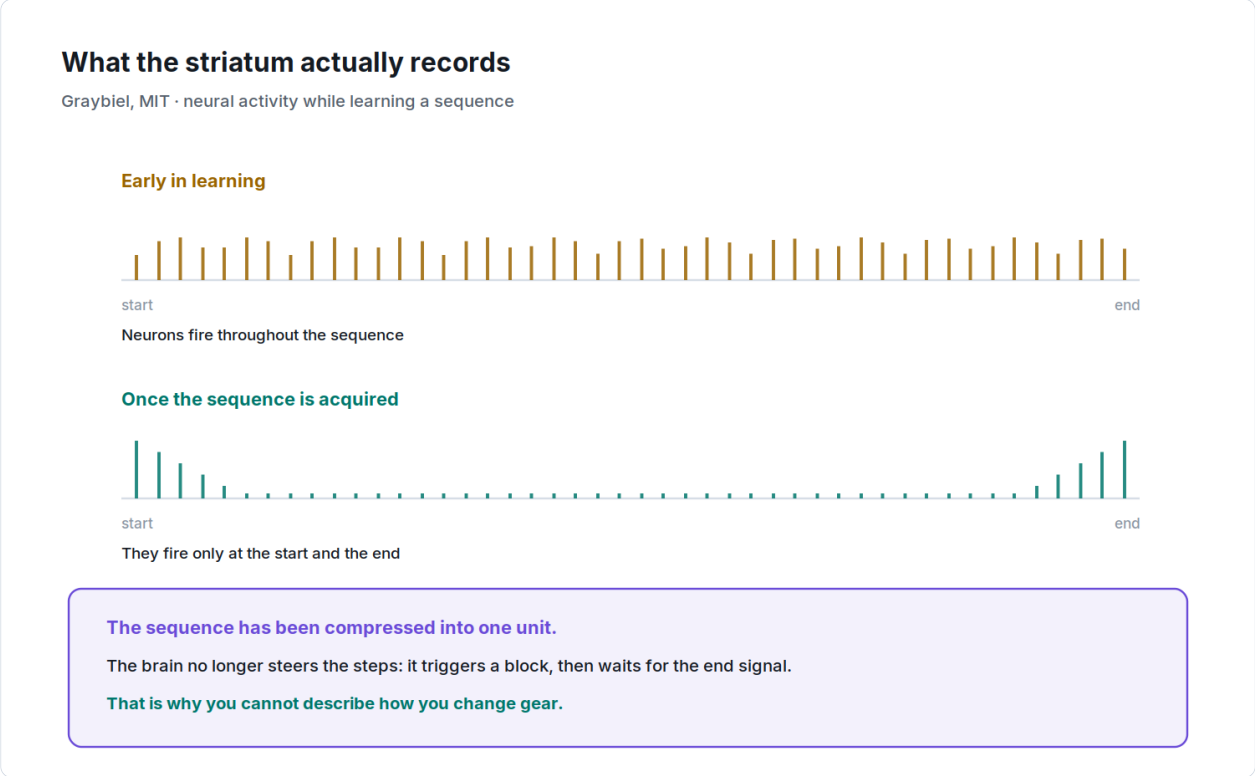
Your teams have already automated how they work. Every workaround, every "we've always done it this way" is a sequence installed in their striatum, stable and free to execute.

When you introduce a standard, you are not filling a void. **You are competing with an automatism that works**, and that holds the crushing advantage of costing nothing.

What the striatum actually records

Ann Graybiel's work at MIT revealed something striking about how these sequences are encoded.

When an animal learns a sequence of actions, the activity of striatal neurons transforms. At first they fire throughout the sequence. Once it is mastered, they fire only at the **beginning** and the **end**, falling silent in between.



Neural activity across the learning of a sequence

The whole sequence has been compressed into a single unit. The brain no longer steers the steps: it triggers a block, then waits for the end signal. This is called chunking.

That is why you cannot describe how you change gear. There are no steps left to describe, only a block whose contents you no longer know.

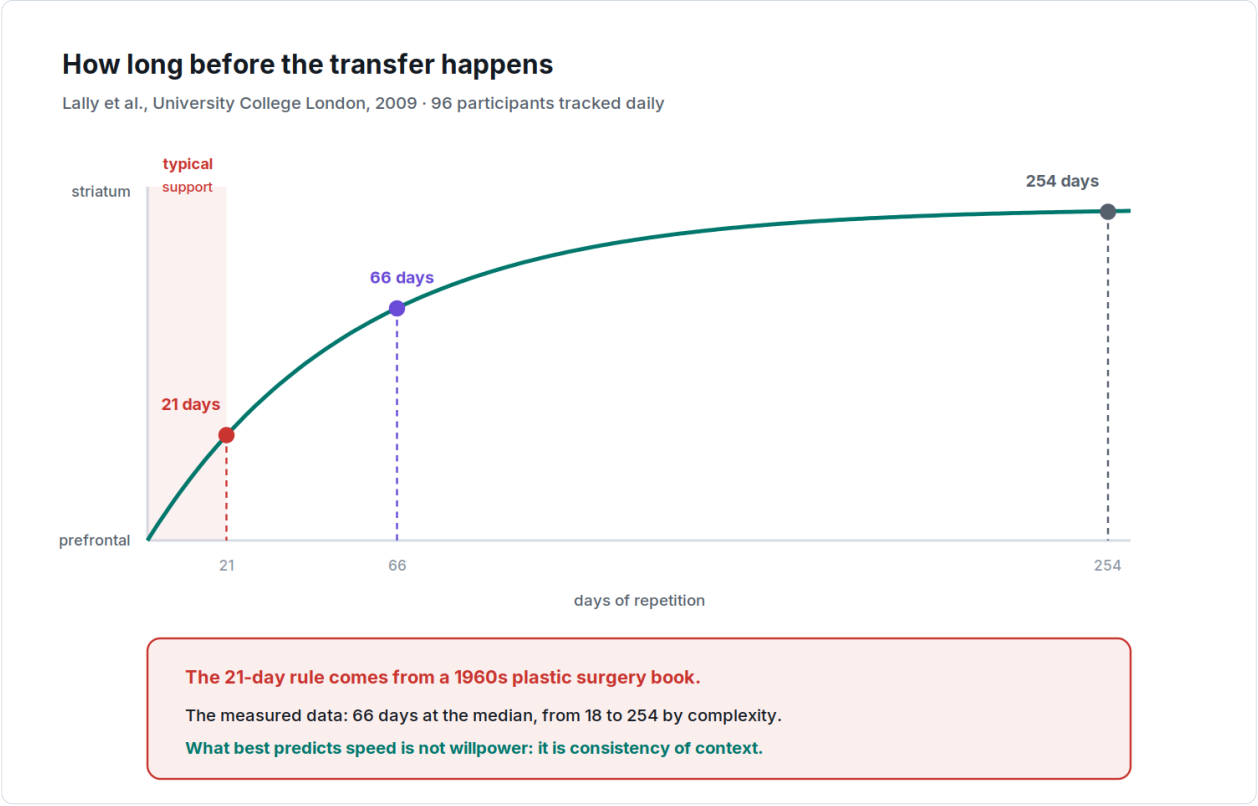
And it is why a standard filed in a binder never becomes a practice. **A binder is not a start signal.** It triggers nothing; it waits to be fetched, which means calling on the expensive circuit.

How long

You've heard it takes twenty-one days to form a habit. That figure rests on nothing: it originates in a 1960s plastic surgery book, where a surgeon observed that his patients took about three weeks to get used to their new face.

The serious data came later.

In 2009, Phillippa Lally and her team at University College London followed 96 participants establishing a new daily behaviour, each reporting every day how automatic the action felt.



The automation curve measured by Lally and her team

Sixty-six days at the median. With considerable spread: from 18 to 254 days depending on the complexity of the action.

Two consequences for anyone running a transformation.

The first is a matter of scale. If your project plans three weeks of support then withdrawal, you leave the floor before the transfer has occurred. The practice is still entirely carried by the prefrontal cortex, which is to say at its most fragile.

The second is more useful. The variable best correlated with speed of installation was neither motivation nor number of repetitions, but **consistency of context**: performing the action at the same moment, in the same place, after the same preceding action.

Which matches chunking exactly. The striatum associates a signal with a sequence. The more stable the signal, the faster the association forms.

The old one doesn't erase

One last point, which explains something every manager has observed.

A well-installed sequence doesn't disappear when you stop using it. It goes quiet, but stays available, and can reactivate if its signal reappears.

You don't remove an automatism. You install another one, which prevails as long as its signal is present.

That is why a standard held for eight months can collapse in a week of pressure. Under strain, the prefrontal cortex, which carried the new movement, is the first to give. The old sequence never needed it.

What to take away

Two structures, two costs. The prefrontal cortex deliberates and tires. The striatum executes and costs almost nothing.

The brain automates to economise. 20% of the energy for 2% of the mass: anything that can shift does.

Chunking **compresses the sequence into a block.** Hence the impossibility of describing a mastered movement, and the importance of the signal that triggers it.

Sixty-six days at the median, up to 254. Three weeks of support installs nothing.

Consistency of context beats willpower. Same moment, same place, same sequence.

Sources

Lally, P., van Jaarsveld, C. H. M., Potts, H. W. W., & Wardle, J. (2009). *How are habits formed: Modelling habit formation in the real world*. European Journal of Social Psychology.

Graybiel, A. M. *Habits, rituals, and the evaluative brain*. Annual Review of Neuroscience. MIT work on chunking and striatal activity across learning.

On cerebral energy consumption: usual estimates place the brain at around 20% of resting expenditure for about 2% of body mass.

CHAPTER 03

Why your standards fade

Four parameters, four structures

8 min read

Here are two standards.

The first has held for eleven years. Nobody reminds anyone of it, nobody audits it, and it is applied by people who have never read it: they learned by watching.

The second was written eight months ago by a competent engineer, validated, presented, posted in the break room. It is applied by one person in three, on bad days.

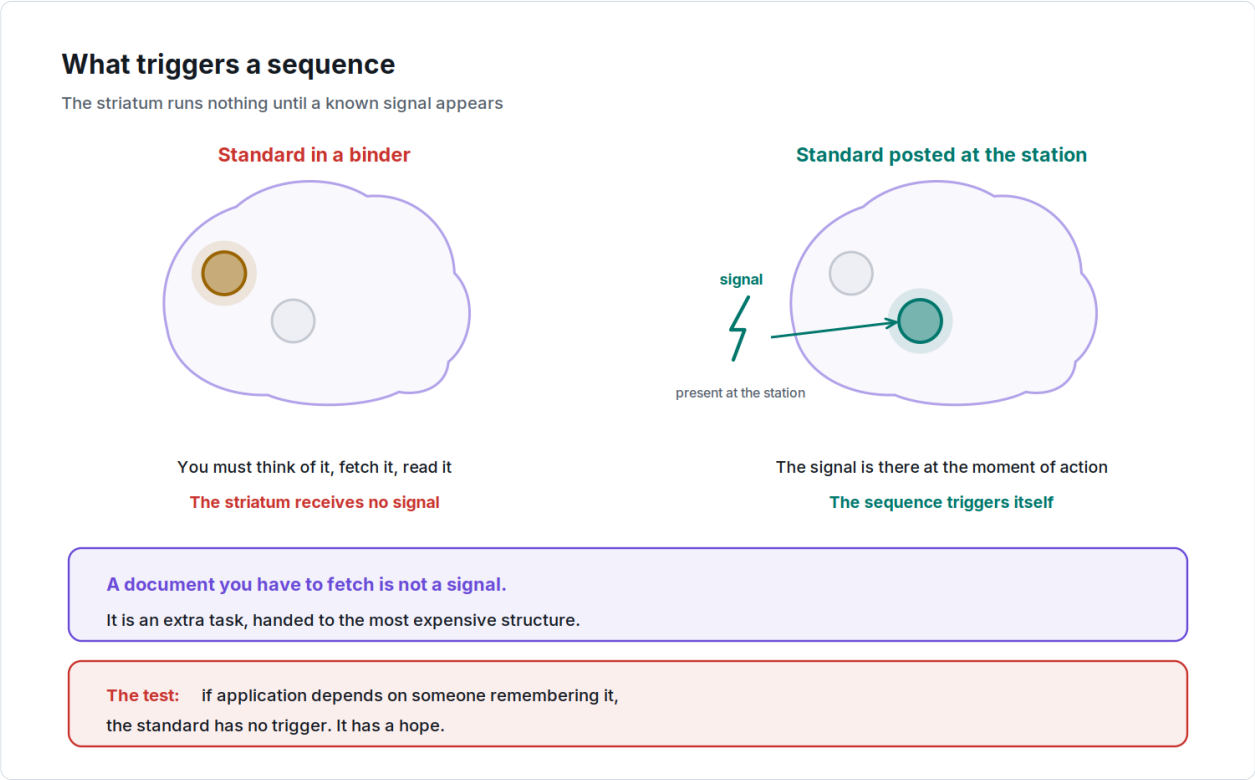
The second is technically superior.

The quality of a standard is not what decides whether it holds.

The signal, first

We saw in the previous chapter that the striatum runs a sequence only if a known signal appears. Not through memory, not through willpower: through a signal.

Most standards have none.



What triggers a sequence, and what triggers nothing

A standard filed in a binder requires, in order to be applied, that you think of it. That means calling on the prefrontal cortex, the structure that costs, at the exact moment you are doing something else.

This isn't viable. On the first day of heavy load, nobody will think of it.

An effective signal has three properties. It is **present at the moment of the action**, not in an office. It is **salient**: it registers without being looked for. It is **constant**: always the same, always in the same place.

> *If applying the standard depends on someone remembering it, the standard has no trigger. It has a hope.*

Four parameters

The trigger is only the first. Three others decide, and each corresponds to a different structure.



The four parameters and the structures that carry them

Load belongs to the prefrontal cortex, whose holding capacity is severely limited, as the next chapter shows. A standard requiring seven points exceeds what the system can carry. It will not be poorly applied out of bad will: it will be **partially** applied, and the forgotten points will differ from one time to the next.

If you audit a standard and find operators each applying seventy per cent of it, but not the same seventy per cent, you have a load problem, not a discipline problem.

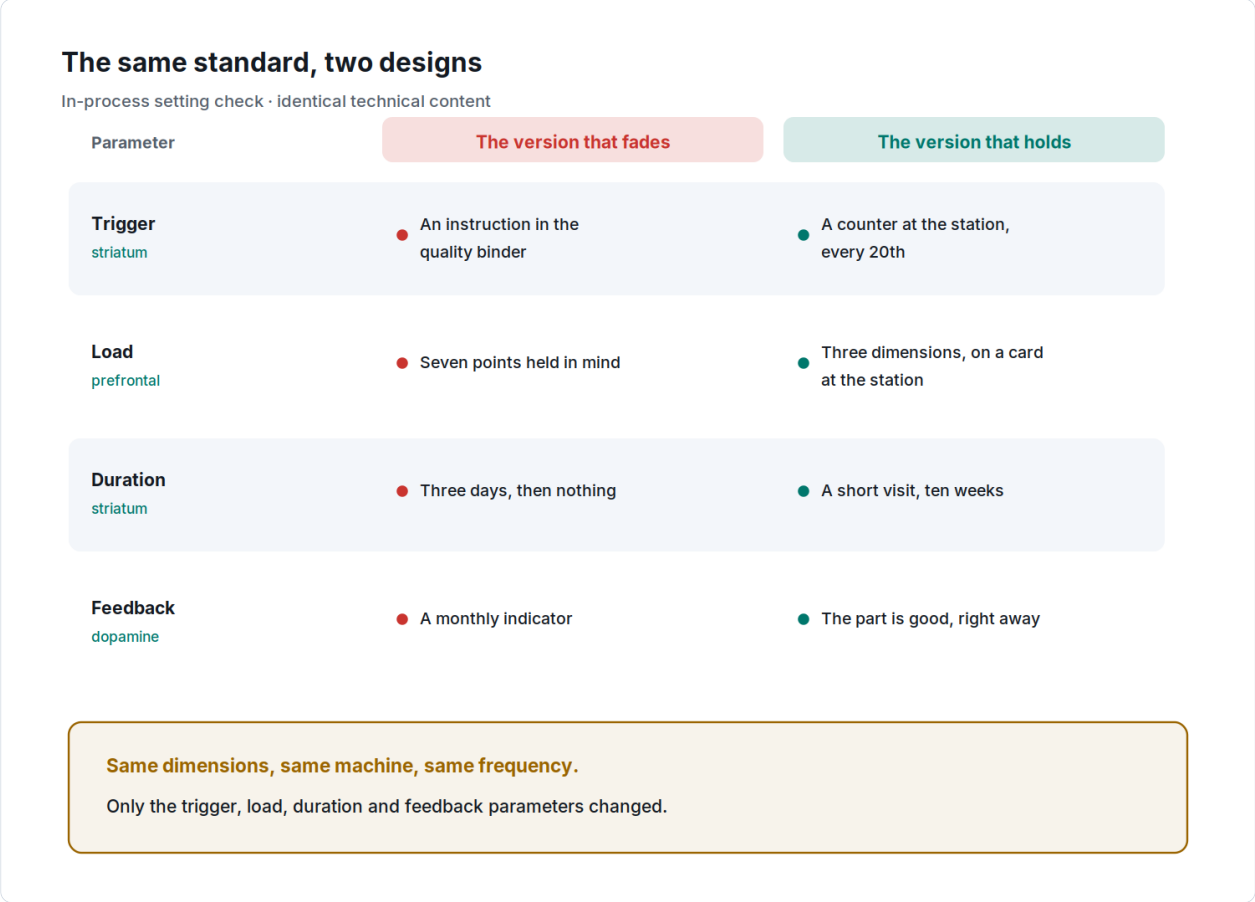
Duration belongs to the striatum and to the time the transfer takes. Sixty-six days at the median. Three weeks of support leaves the practice entirely carried by the structure that tires.

Feedback belongs to dopamine, which chapter 8 shows encodes not pleasure but the gap against prediction. A standard whose benefit appears in a monthly indicator reinforces nothing: the link between the action and its effect cannot be established.

These four parameters work **in series**. One missing is enough to bring the whole down.

The same standard, twice

Take a concrete case and run it through all four.



An in-process check, designed two ways

Note what doesn't change: **the technical content is rigorously identical**. Same dimensions, same machine, same reason, same frequency.

What changes belongs entirely to the design of the execution conditions. Nothing requiring trade competence, nothing expensive, nothing needing approval.

One of the two versions will still be applied in two years. The other will be gone by Christmas.

Audit the standard, not the people

When a standard isn't applied, the usual approach is to question people: who applies it, who doesn't, why.

That approach always produces the same answers, "we don't have time", "we forgot", and the same actions: reminders, awareness sessions, tighter controls.

None of them corrects any of the four parameters. Hence the three-week rebound, then the return to baseline.

Four questions do better, and they bear on an object rather than on people.

Is there a signal, present at the workstation, at the moment of the action?

How many elements must be held in mind to apply it?

How long ago did the support stop?

What visibly happens in the seconds that follow?

None of these four questions puts anyone at fault. They describe a device you can modify without asking anyone to change.

What to take away

Technical content is not the subject. Two identical versions of a standard meet opposite fates depending on their execution conditions.

The signal is most often absent. A document you have to go and fetch isn't one.

Four parameters, four structures. Trigger and duration for the striatum, load for the prefrontal cortex, feedback for dopamine.

They work in series. One missing link brings the whole down.

The signature of an overloaded standard: everyone applies part of it, and not the same part.

Sources

The mechanisms of automation, triggering and duration are detailed in chapter 2, with their sources, notably Lally et al. (2009) and Graybiel's work on sequence encoding.

Working memory load is covered in chapter 4; dopaminergic reinforcement in chapter 8.

CHAPTER 04

Four slots, no more

Why visual management actually works

8 min read

"We put colours everywhere, so what?"

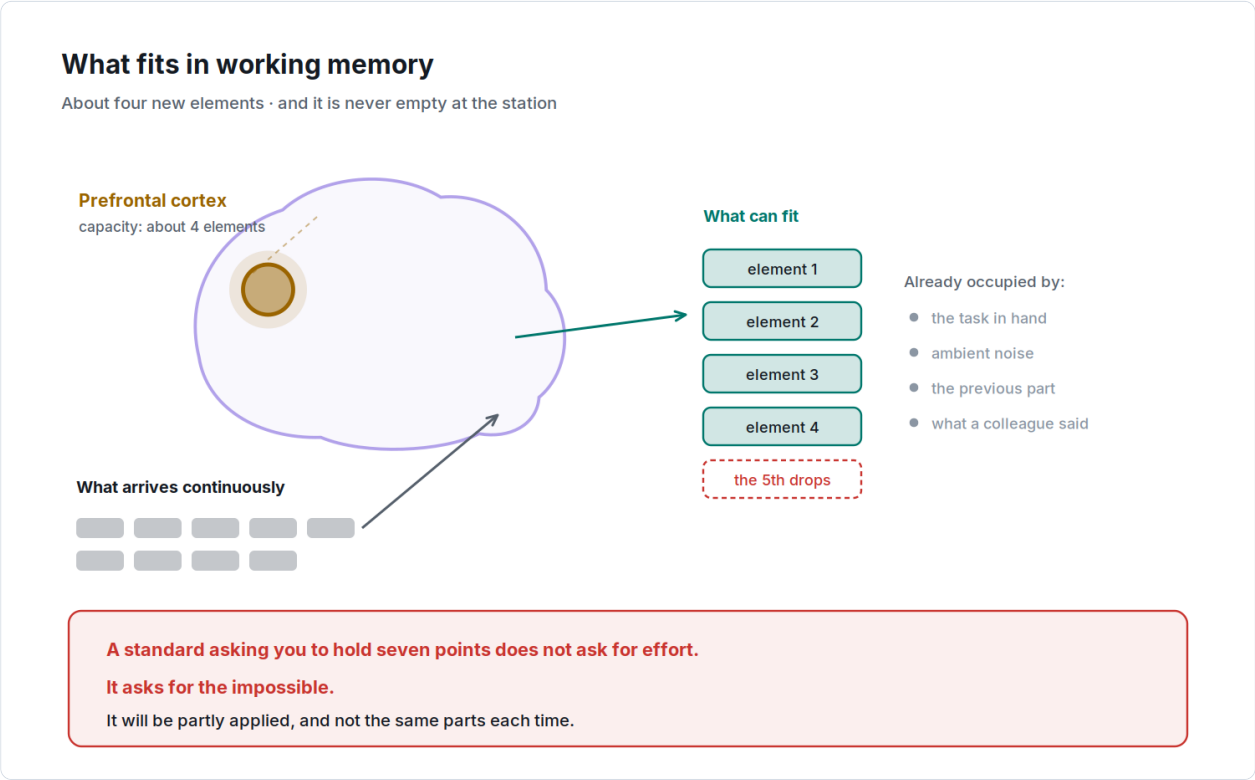
That sentence, said in every workshop by someone who has watched three initiatives come and go, is often right. A great deal of visual management **is** decorative.

But when it's well designed, it acts on the scarcest resource in human work, and we can say precisely which one.

The real capacity

The prefrontal cortex supports what is called working memory: the space where you hold, temporarily, the information you are manipulating. The number you keep in mind long enough to dial it. The instruction you hold long enough to apply it.

That space is narrow.



What working memory can hold

George Miller proposed in 1956 a capacity of about seven elements. More recent estimates, notably Nelson Cowan's, are harsher: **around four** for new material that cannot be grouped.

Four. And more to the point: that space is never empty while you work. It is already occupied by the task in hand, by noise, by the previous part that caused trouble, by what a colleague just asked.

A standard requiring seven points held in mind doesn't demand an effort. **It demands the impossible.**

Three loads, only one you can change

In 1988, John Sweller published in *Cognitive Science* the work that became cognitive load theory. He distinguishes three sources of occupation for that space.

Intrinsic load comes from the complexity of the task itself. You can reduce it by breaking the task down, never remove it.

Germane load serves learning: building the representation, understanding why this action rather than another. It is the only one you want to maximise.

Extraneous load comes from how the information is presented. It adds nothing and consumes resource.

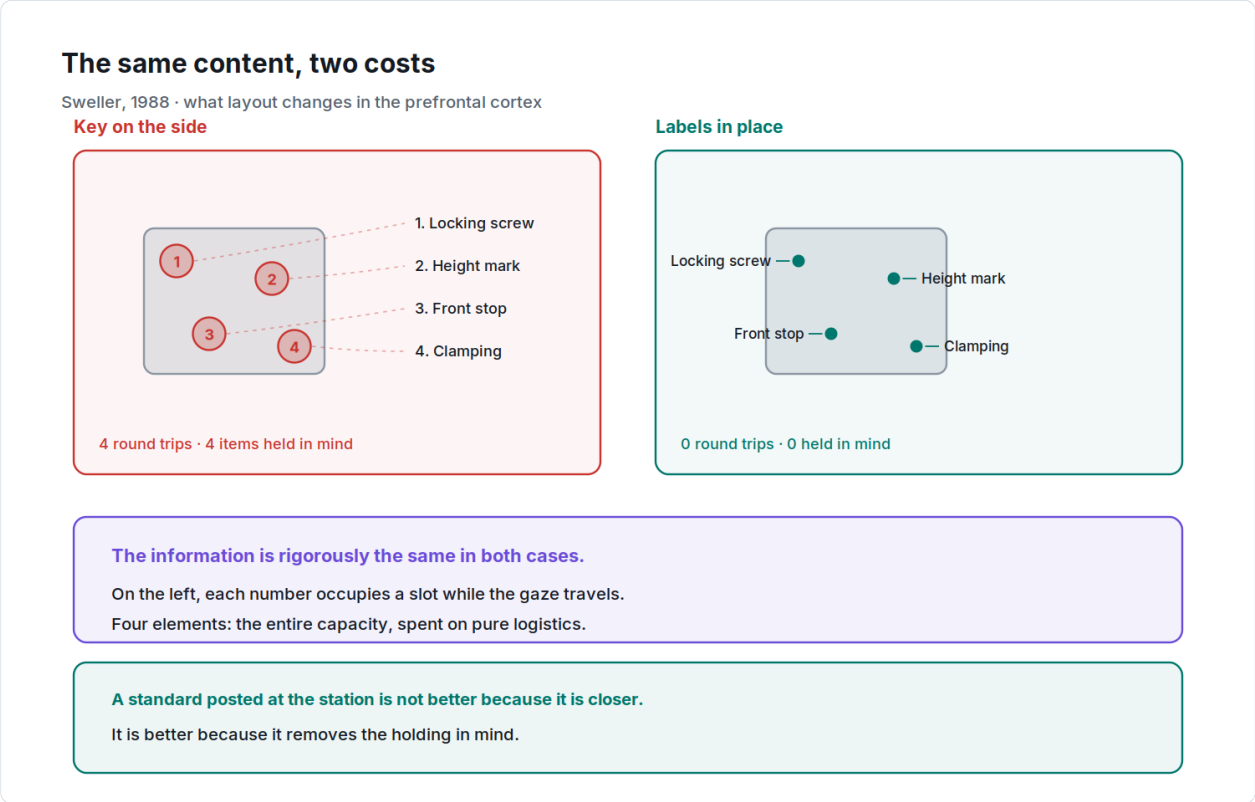
The third is the one that matters, because it is **entirely under your control**. It depends neither on the trade nor on people's capacities. It depends on layout.

The experiment that proves it

Sweller and colleagues demonstrated an effect whose application to the workstation is direct.

Take a technical diagram with its key on the side: numbered markers on the drawing, the list of correspondences to the right. To understand marker 3 you must see it, **memorise "3"**, move your gaze, find the line, read it, then return holding what you just read in mind.

Now take the same diagram with the labels placed beside each element.



Two layouts of the same information

The information is rigorously identical. But in the first case, every element consumes a slot of working memory during the journey of the gaze, for an operation that teaches nothing.

Four elements, four slots. **The entire capacity, spent on pure logistics.**

This is called the split-attention effect. It is one of the most robust findings in the theory.

Its consequence explains something every practitioner has noticed: a standard posted at the workstation isn't better than a standard in a binder because it's closer. It's better because it **removes the holding in mind**.

The binder forces you to remember. The posting lets you look.

Three rules that follow

The information must be where the action happens. Every metre walked, every click, every "I'll look later" is extraneous load that will eventually win.

What matters must stand apart from what doesn't. A board showing twenty indicators shows nothing: working memory will take four, and not necessarily the right ones.

The gap must be visible without being calculated. This is the point most often missed.

A board showing "Output today: 847" demands three mental operations: compare to target, assess the gap, judge whether it's normal. Three slots out of four.

A board showing a green or red cell demands none. The gap is already calculated, already judged.

> *That is the difference between showing data and showing a state. The first leaves the work to the reader; the second has already done it.*

A limit worth stating

Sweller's work was done in a **learning** context, with learners in acquisition. The split-attention effect is solidly documented there.

An experienced operator isn't in that situation: their intrinsic load is far lower, since much of the action has passed to the striatum.

One finding is worth flagging: a support device that helps a beginner can **hinder** an expert, by forcing them to process information they already hold.

This doesn't condemn visual management. It suggests separating what serves learning from what serves execution, and not imposing the first on those who no longer need it.

What to take away

About four elements. And the space is never empty in a work situation.

Only one load concerns us. Extraneous load depends solely on how information is laid out.

The split-attention effect is the central mechanism. Any information you must go and find occupies a slot during the journey.

Showing a state beats showing data. The first has already done the interpretive work.

A support that helps the beginner can hinder the expert. Their intrinsic load is not the same.

Sources

Sweller, J. (1988). *Cognitive load during problem solving: Effects on learning*. Cognitive Science, 12(2), 257–285.

Sweller, J., van Merriënboer, J. J. G., & Paas, F. (1998). *Cognitive architecture and instructional design*. Educational Psychology Review.

Miller, G. A. (1956). *The magical number seven, plus or minus two*. Psychological Review, 63(2), 81–97.

Cowan, N. (2014). *Working memory underpins cognitive development, learning, and education*. Educational Psychology Review.

Chandler, P. & Sweller, J. (1991). Work on the split-attention effect.

CHAPTER 05

The brain knows before you do

Fifty milliseconds, and nobody says a word

9 min read

You type a password. Your fingers get it wrong. And before you've even read the screen, before any formed thought, you know it failed: you're already correcting.

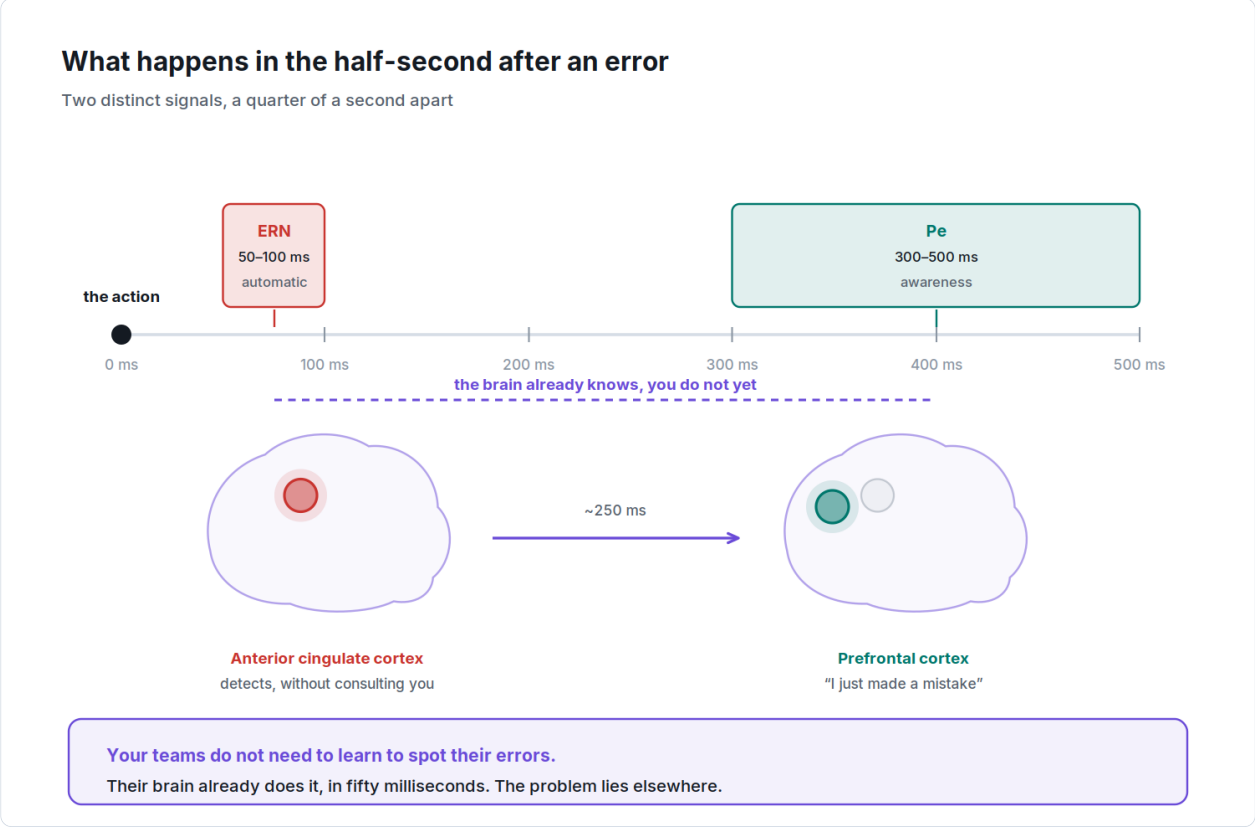
That sensation has a measurable electrical correlate. It's called the ERN, for error-related negativity.

It appears **fifty to a hundred milliseconds after the faulty action**.

Two signals, a quarter of a second apart

Since the 1990s we've known how to record this signal at the surface of the skull. It appears whenever someone makes an error in a fast task, and its main source lies in the anterior cingulate cortex.

But the ERN isn't alone. A second signal follows, very different.



The timeline of error detection

The Pe, for error positivity, appears **three hundred to five hundred milliseconds** after the action. And it is this one, not the ERN, that is associated with becoming aware of the error.

In other words: your brain detected the fault and began correcting for a quarter of a second **before** you knew you were wrong.

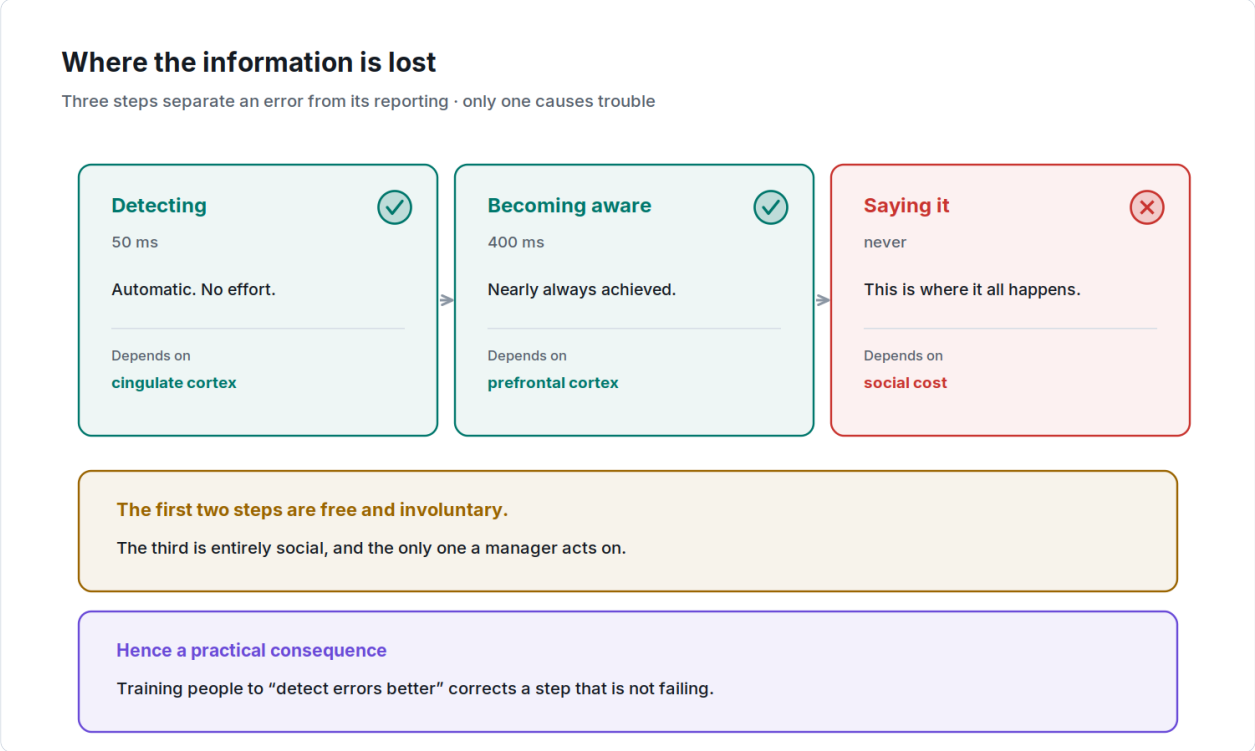
The clearest demonstration comes from patients with lesions of the cingulate cortex. In them, the electrical distinction between error and correct response is sharply reduced: fast detection is impaired.

And yet they remain able to **say** they made a mistake. Awareness of the error survives the destruction of the detector.

What they lose is something else: the ability to **improve their accuracy after an error**. The fast signal isn't there to know. It's there to adjust.

What this moves

This result completely changes how the problem of speaking up should be posed.



The three steps between an error and reporting it

Detecting is **automatic, free, involuntary**. Nobody needs to learn it.

Becoming aware arrives a quarter of a second later, nearly always.

Saying is the only step that causes trouble, and it depends on no detection structure. It depends on a social calculation.

The consequence is direct: **training people to "spot their errors better" corrects a step that isn't failing**. The information already exists, in fifty milliseconds, in every operator's brain. It just doesn't come out.

Why errors hurt

A second finding illuminates why the saying step is so costly.

The ERN isn't a neutral signal. Its amplitude varies with the **importance** attached to the error: it rises when instructions stress accuracy, it rises in anxious individuals.

More troubling still: work has shown that errors trigger a defensive response, measurable physiologically. The authors conclude that **errors are aversive**, in the literal sense. They aren't merely a finding, they are unpleasant.

The cingulate cortex is, moreover, involved in processing pain and negative affect, alongside cognitive control.

What this suggests: at the moment someone considers reporting, they have already undergone an unpleasant reaction. Asking them on top of that to expose the error publicly is adding a social cost to one already paid.

What organisational research found, without knowing it

In 1991, Amy Edmondson, a doctoral student at Harvard, crossed medication-error data with measures of teamwork quality in hospital wards.

Her hypothesis was obvious: better teams must make fewer errors.

The result was the reverse. **The wards with the best teamwork reported the most errors.**

After weeks looking for her methodological mistake, she formulated the other reading: these teams weren't making more errors, they were able to talk about them. What had been measured wasn't the number of errors **made**, but the number **reported**.

Put that result next to the ERN, and it becomes mechanically obvious.

The number of errors made is roughly constant from one ward to another: the human brain doesn't vary that much. Detection is constant too: fifty milliseconds everywhere.

What varies from ward to ward is only the third step.

A ward with no reported incidents isn't a safe ward. It's a ward where the third step isn't crossed.

What sets the cost

Since the problem lies entirely there, let's look at what determines it.

The immediate reaction. What happens in the thirty seconds after a report is watched by the whole team. Not the stated policy: the actual reaction. A sigh, a "how did that happen?" in a certain tone, or a problem fixed without a word for the person who raised it. The price has just been set.

What happens to the information. A report that leads to nothing teaches that speaking is pointless. The cost was real, the benefit nil. Repeat that a few times and the behaviour extinguishes.

What happens when the person is wrong. This is the decisive test. An operator stops the line, checks are made, everything is in order, twenty minutes lost. If the message received is "you were right to check", reporting stays possible. If it's "next time, be sure", nobody will report before they are certain, which is to say too late.

> *A team doesn't learn from what it's told about the right to be wrong. It learns from what happens to the first person who exercises it.*

A limit worth stating

Two cautions on this chapter.

The first concerns the ERN itself. Its exact source is debated: most work places it in the dorsal anterior cingulate cortex, but recent analyses propose a more posterior origin. That doesn't change the phenomenon, a fast signal exists, but the precise localisation isn't settled.

The second matters more. The ERN is measured in the laboratory, on tasks of a few hundred milliseconds, with consequence-free errors. An operator fitting a part backwards isn't in that situation.

I assume the fast detection mechanism also operates at the workstation. That's plausible, it's a general action-monitoring mechanism, but it isn't demonstrated in that context.

What holds without the extrapolation: error detection is early and largely automatic, and the organisational problem sits downstream.

What to take away

The ERN appears in 50 to 100 ms. The Pe, associated with awareness, in 300 to 500 ms. Your brain knows before you do.

The detector isn't there to know, but to adjust. Patients with lesions know they were wrong, but no longer correct their accuracy afterwards.

Errors are aversive in the literal sense. They trigger a measurable defensive response.

Three steps, only one failing. Detecting and becoming aware are automatic. Saying is social.

Silence doesn't measure safety. It measures what it costs to speak.

Sources

Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993). Original description of the ERN.

Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1991). Independent discovery of the same signal.

Hajcak, G. & Foti, D. (2008). *Errors are aversive: defensive motivation and the error-related negativity*. Psychological Science.

On patients with cingulate lesions: work showing impaired fast detection with preserved awareness of the error, and loss of post-error adjustments.

Edmondson, A. C. (1999). *Psychological safety and learning behavior in work teams*. Administrative Science Quarterly, 44(2), 350–383.

On the debate about the ERN's source: localisation analyses largely place the generator in the dorsal anterior cingulate cortex, with a posterior origin under discussion.

CHAPTER 06

Under pressure, the brain changes pilot

Why a stressed organisation turns rigid

9 min read

A team put under pressure to improve becomes less able to improve.

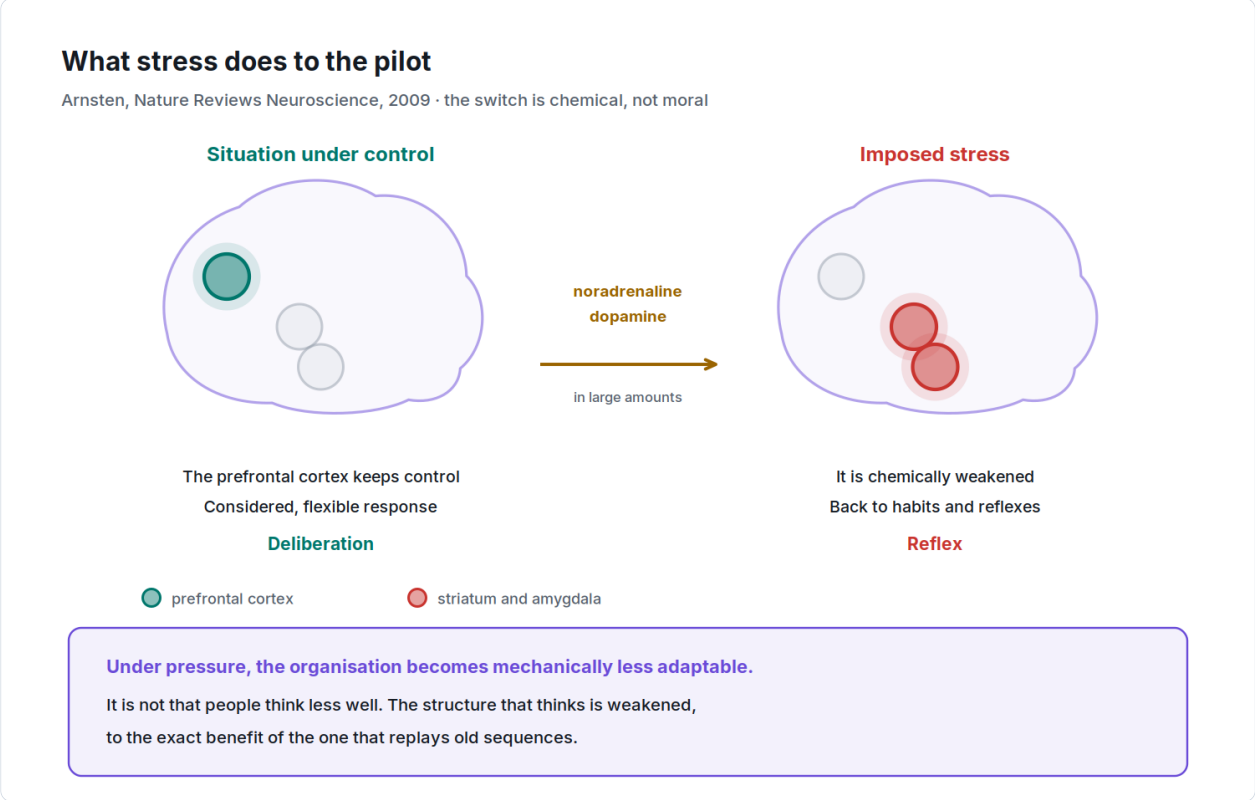
This isn't a management paradox. It follows directly from what happens in the brain when pressure is perceived as imposed.

The switch

The prefrontal cortex is, of all brain regions, the most sensitive to stress.

Amy Arnsten's work, synthesised in *Nature Reviews Neuroscience*, describes a precise mechanism. Under normal conditions, moderate amounts of noradrenaline and dopamine optimise its functioning.

Under stress, those same neuromodulators are released in **large amounts**, and the effect reverses.



What stress does to control of action

The prefrontal cortex is rapidly weakened. And in the same movement, the responses carried by the striatum and the amygdala are **strengthened**.

This isn't a general drop in performance. It is a **transfer of control**: the structure that deliberates hands over to those that replay learned sequences and evaluate threat.

The mechanism has obvious survival value. Facing immediate danger, a reflex beats deliberation. It becomes costly when the situation demands, on the contrary, finding a new solution.

You know this scene

You know this sequence, and chapter 2 gave you half of it.

A team has installed a new standard. It has held for six weeks, still carried by the prefrontal cortex: the transfer to the striatum isn't complete.

Then a surge in load. An urgent customer, a breakdown, an absence.

What we already knew: the expensive system gives way to the free one.

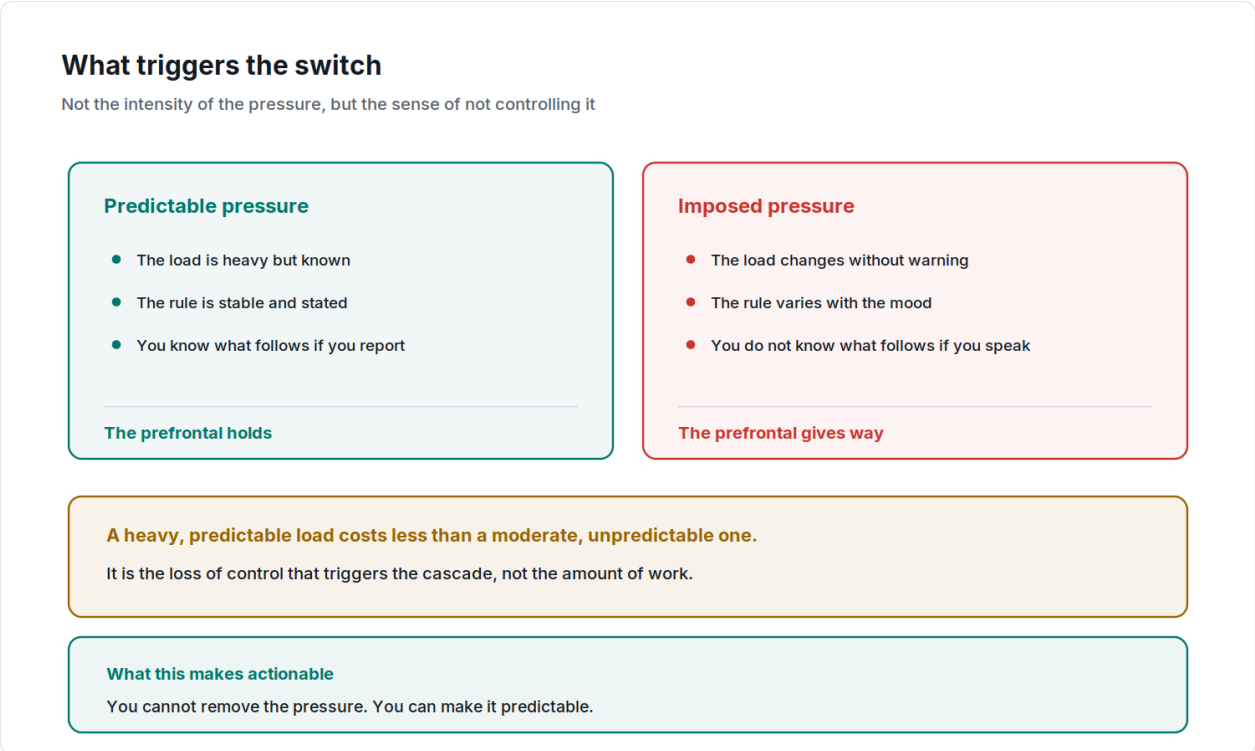
What we now know: stress actively accelerates that switch. It doesn't merely tire the prefrontal cortex. It weakens it chemically while strengthening the striatum, which is to say strengthening precisely the old sequence you were trying to replace.

The organisation gets the exact opposite of what it is asking for.

The variable that decides

The most useful point in this work isn't that stress degrades the prefrontal cortex. It's **which kind** of stress does.

Studies converge on a precise criterion: what triggers the cascade is the sense of **not controlling** the situation. Even mild stress, if imposed, is enough. Uncontrollability is what counts, not intensity.



What triggers the switch

That changes everything for a manager, because workload is often not negotiable, whereas its predictability is.

A heavy but predictable load costs less than a moderate but unpredictable one. A strict but known rule costs less than a lenient but variable one.

It also explains something documented at organisational level long before the mechanism was known: under perceived threat, groups turn rigid. They fall back on familiar behaviours, restrict the information that circulates, tighten control. The effect was described as early as 1981.

A group of brains whose prefrontal cortices are simultaneously weakened produces exactly that behaviour.

What this says about handling errors

Back to the previous chapter. We left the problem at the third step: detecting is free, becoming aware is too, **saying** is costly.

The stress mechanism now lets us say why some organisations never manage to bring that cost down.

When the response to a report is **unpredictable**, lenient one day, punished the next, depending on mood or on what's at stake, every report becomes an uncontrollable situation. The brain of the person considering speaking up is in exactly the configuration described above: weakened prefrontal cortex, strengthened amygdala.

They will not deliberate. They will do what a system under threat does: take the safest response, which is silence.

So it isn't the absence of sanction that protects reporting. It is the predictability of the response.

Why "no blame" doesn't work

This illuminates a very widespread piece of advice, and shows why it fails.

You read everywhere that all sanction should be removed to free up speech. James Reason, one of the founders of the study of organisational error, had written plainly that a general amnesty on all dangerous acts would lack credibility and appear contrary to elementary justice.

The mechanism explains why.

Undifferentiated leniency is not predictable. If nothing is ever sanctioned, nobody knows where the line runs, and an invisible line is exactly an uncontrollable situation.

A conscientious operator who watches a colleague bypass a safety device without consequence doesn't draw safety from it. They draw the conclusion that the rule has no value, so they no longer know what is expected.

> *A clear boundary protects better than vague leniency. In the first case, everyone knows where they stand.*

What Reason proposes instead, what is called just culture, rests on a distinction known in advance: honest error calls for correcting the system, at-risk behaviour calls for understanding why the shortcut was rational, reckless behaviour calls for sanction.

What matters isn't the level of leniency. It's that the rule be **stated beforehand** and **held consistently**. That is what makes the situation controllable.

A limit worth stating

Two cautions.

On the extrapolation. Arnsten's work is mainly on primates and rodents, with laboratory stressors: intense white noise, brief restraint. The link with a durably oppressive work climate is plausible but not demonstrated as such.

On a shortcut to avoid. You often read that threat "shuts down the prefrontal cortex" or produces an "amygdala hijack". Those formulas are too clean. What is documented is a relative weakening of prefrontal control and a relative strengthening of subcortical responses, not an extinction.

The organisational phenomenon stands independently: under threat, groups turn rigid. It doesn't need a neurobiological veneer to hold.

What to take away

Stress transfers control. The prefrontal cortex is weakened, the striatum and amygdala strengthened.

The variable is controllability, not intensity. Mild but imposed stress is enough; a heavy but predictable load costs less.

An organisation under pressure becomes mechanically less adaptable. It replays its old sequences at the moment it most needs new ones.

Silence calls for more predictability, not more kindness. A variable response makes every report uncontrollable.

A clear boundary beats vague leniency. It's consistency, not the level of sanction, that protects speech.

Sources

Arnsten, A. F. T. (2009). *Stress signalling pathways that impair prefrontal cortex structure and function*. Nature Reviews Neuroscience, 10(6), 410–422.

Arnsten, A. F. T. & Goldman-Rakic, P. S. (1998). *Noise stress impairs prefrontal cortical cognitive function in monkeys*. Archives of General Psychiatry, 55(4), 362–368.

Liston, C., McEwen, B. S., & Casey, B. J. (2009). *Psychosocial stress reversibly disrupts prefrontal processing and attentional control*. PNAS, 106(3), 912–917.

Staw, B. M., Sandelands, L. E., & Dutton, J. E. (1981). *Threat-rigidity effects in organizational behavior*. Administrative Science Quarterly, 26(4), 501–524.

Reason, J. (1997). *Managing the Risks of Organizational Accidents*. Ashgate. On just culture and the impossibility of a blame-free culture.

CHAPTER 07

The work happens during the breaks

What the hippocampus replays while you do nothing

8 min read

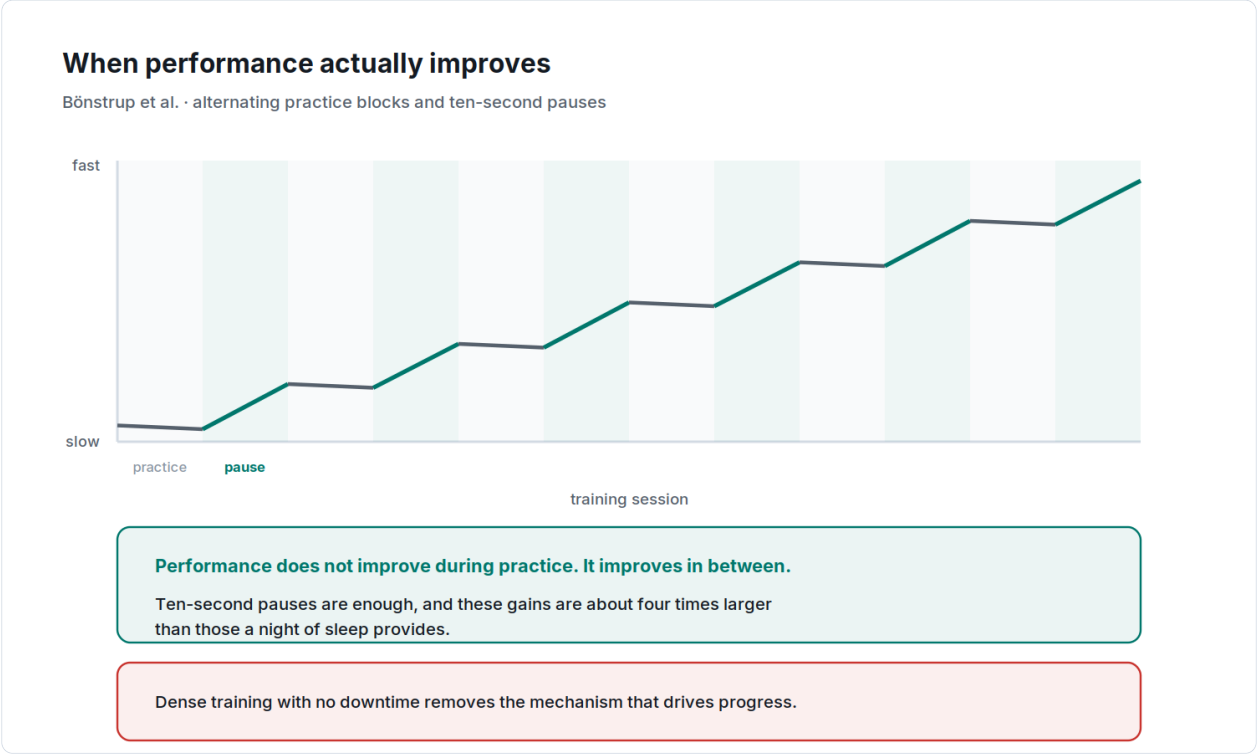
Three days of intensive training, Monday to Wednesday, with no downtime.

That's the standard format. It's also the worst one you could design, for a reason that has nothing to do with participant fatigue.

Where the gains happen

A team of researchers did something simple: measure performance not from one session to the next, but **within** a session, second by second.

Participants learned a keyboard sequence, alternating ten-second blocks of practice with ten-second pauses.



Where performance actually improves within a session

The result contradicts the entire intuition behind training.

Performance barely improves during the practice blocks. It improves during the pauses.

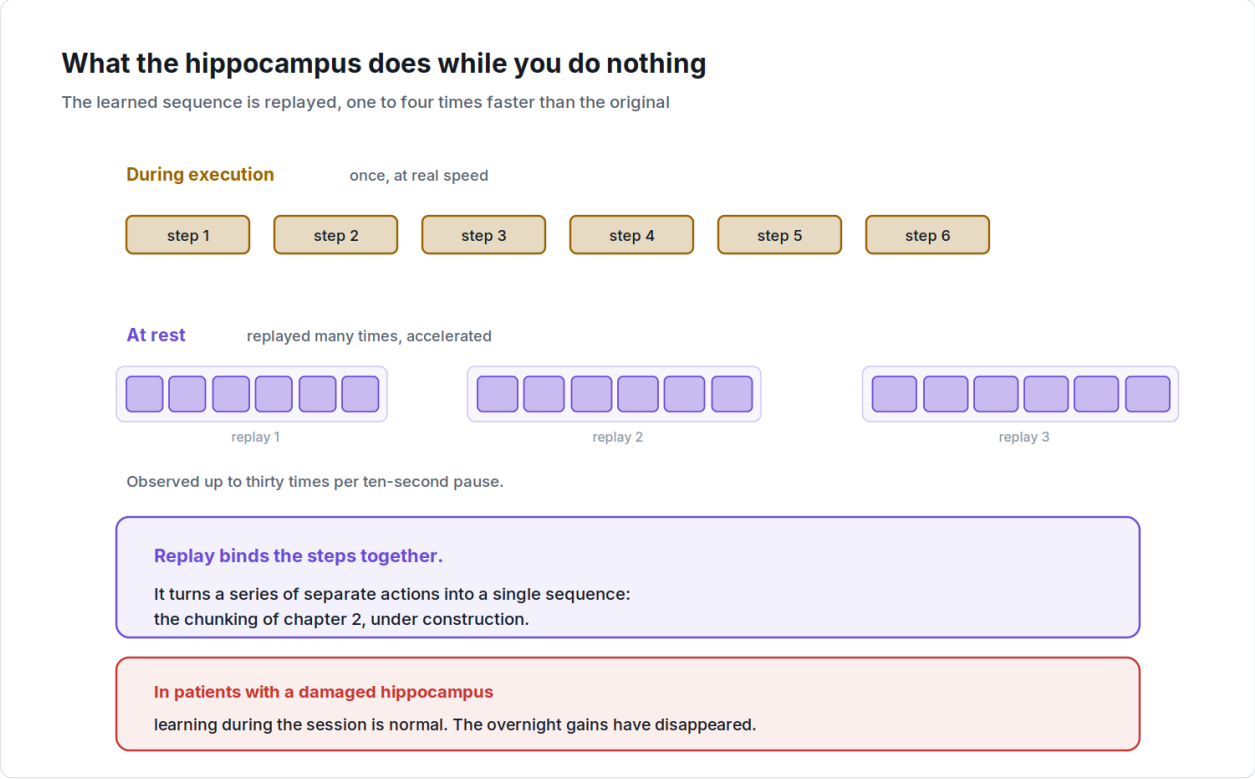
These gains, called micro-offline gains, account for most of the progress in early learning. And their magnitude is about **four times greater** than the gains usually attributed to a night's sleep.

Ten seconds of doing nothing is worth more than ten seconds of additional repetition.

What happens during those ten seconds

The hippocampus doesn't rest. It replays.

Recording brain activity during these short pauses reveals spontaneous reactivation of the neural patterns corresponding to the sequence just practised. This is called replay.



What the hippocampus does during pauses

Two features make this replay remarkable.

It is accelerated. The sequence is replayed one to four times faster than its real execution.

It is massive. As many as thirty events have been detected per ten-second pause.

In other words: while the participant does nothing, their brain replays the sequence dozens of times, at speed.

And that replay does precise work: it **binds the steps together**. This is the chunking of chapter 2 under construction, a series of separate movements becoming a single unit.

Proof by lesion

How do we know the hippocampus is necessary, and not merely present?

Through patients in whom it is damaged.

On a motor sequence task, these patients learn **normally during the session**. Their performance improves like controls'.

But the gains observed after a night, the ones controls obtain without further practice, **have disappeared**.

The hippocampus isn't needed to learn during training. It's needed for learning to consolidate once you stop.

Later work suggests this same structure is also involved in the gains obtained during short waking pauses, not only during sleep.

What this condemns

Take the most widespread training format in industry.

Three consecutive days, dense, with no downtime. The consolidation mechanism has almost no opportunity to work. Participants leave with the impression of having seen a great deal, and minimal consolidation.

A full day of practice with no breaks. The same at smaller scale: repetitions piled into a window where replay cannot operate.

A five-day project followed by nothing. The worst of both: maximum density during, then no opportunity to consolidate after.

Conversely, what the mechanism suggests:

Short, spaced sessions beat one long session of equivalent duration.

Real pauses during practice, where nothing is done: not a different exercise, not a discussion, not the phone.

A night between two sessions, which adds further consolidation.

That last point deserves emphasis: two half-days separated by a night aren't equal to one day. They're better.

A limit worth stating

Three cautions on this chapter.

On the nature of the tasks. The work concerns simple motor sequences, typing a run of keys. An industrial standard is more complex, longer, and performed in a noisy environment. The transposition is plausible; it isn't measured.

On the length of pauses. Ten seconds suffices in the laboratory, on a task lasting a few seconds. Nothing indicates ten seconds suffices after twenty minutes of real work.

On what "pause" means. In these protocols, the pause is genuine inactivity, without stimulation. A pause spent on your phone is probably not equivalent, which is plausible but, as far as I know, not established in this precise context.

What holds without extrapolation: consolidation is an active process that occurs **outside practice**, it depends on the hippocampus, and it can be prevented by depriving people of rest.

What to take away

Gains occur during pauses. In the early phase, most progress happens between practice blocks, not during them.

These micro-offline gains **exceed those of sleep.** By about four times, on available measures.

The hippocampus **replays the sequence at speed**, dozens of times per pause. That replay binds the steps: chunking under construction.

Patients with lesions learn but don't consolidate. The structure isn't needed during, it's needed after.

Dense training removes the mechanism. Short, spaced sessions beat a three-day block.

Sources

Bönstrup, M., Iturrate, I., Thompson, R., Cruciani, G., Censor, N., & Cohen, L. G. (2019). *A rapid form of offline consolidation in skill learning*. Current Biology. Identification of micro-offline gains.

Bönstrup, M. et al. (2020). Large-scale confirmation of gains during short pauses.

Schapiro, A. C. et al. (2019). Absence of overnight gains on a motor sequence task in patients with hippocampal lesions, with learning preserved during training.

Jacobacci, F. et al. (2020). Link between hippocampal activity at rest and the magnitude of gains.

On accelerated replay in slow-wave sleep: intracortical recordings showing reactivation of learned patterns at one to four times execution speed.

On targeted reactivation during sleep and its effects on procedural consolidation: work using auditory or olfactory cues.

CHAPTER 08

The day nothing surprises anymore

Dopamine does not encode pleasure

9 min read

In the early 1990s, in a laboratory in Fribourg, Wolfram Schultz recorded the activity of dopamine neurons in monkeys.

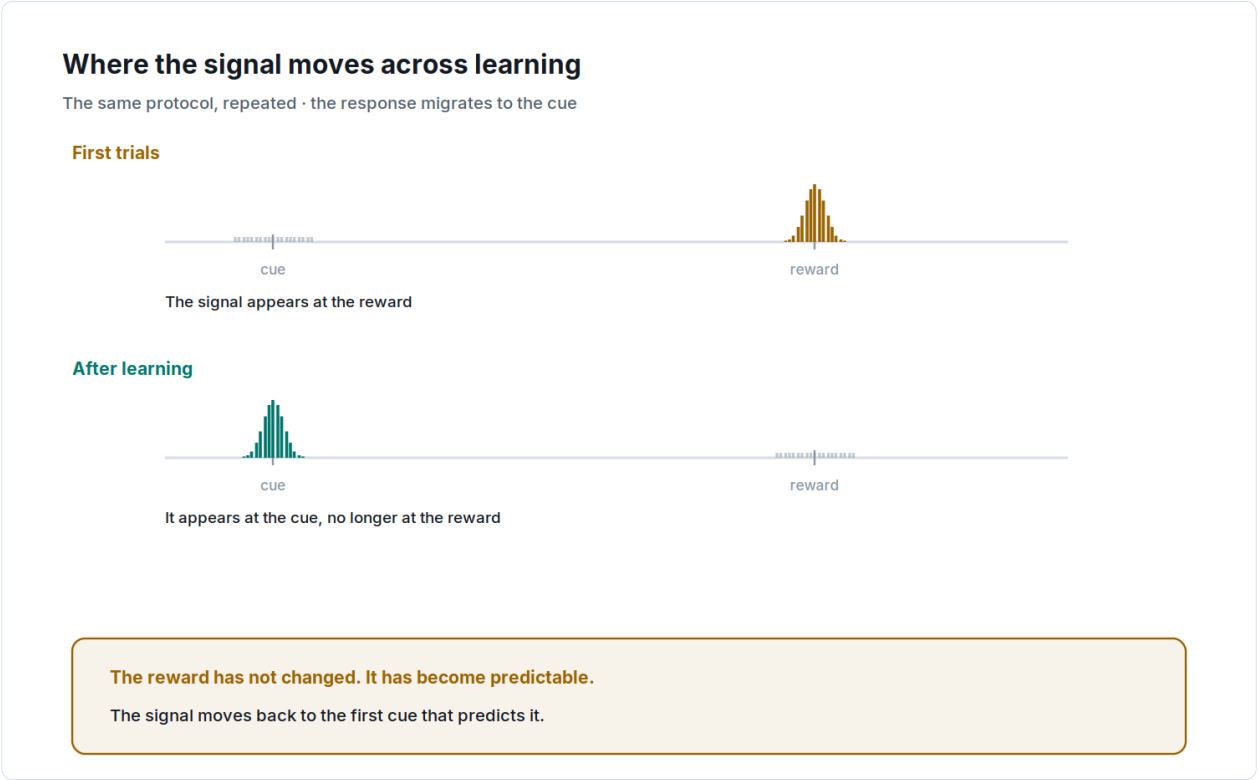
The protocol was simple: the animal receives a drop of juice, you observe the neural response. The established theory was equally simple: dopamine signals pleasure. Juice, firing.

Then Schultz introduced a light that preceded the juice. The animal quickly learned to anticipate it.

And the neural response changed.

The signal moves

Once the association was learned, the neurons no longer responded to the arrival of the juice. **They responded to the light.**



Where the signal moves across learning

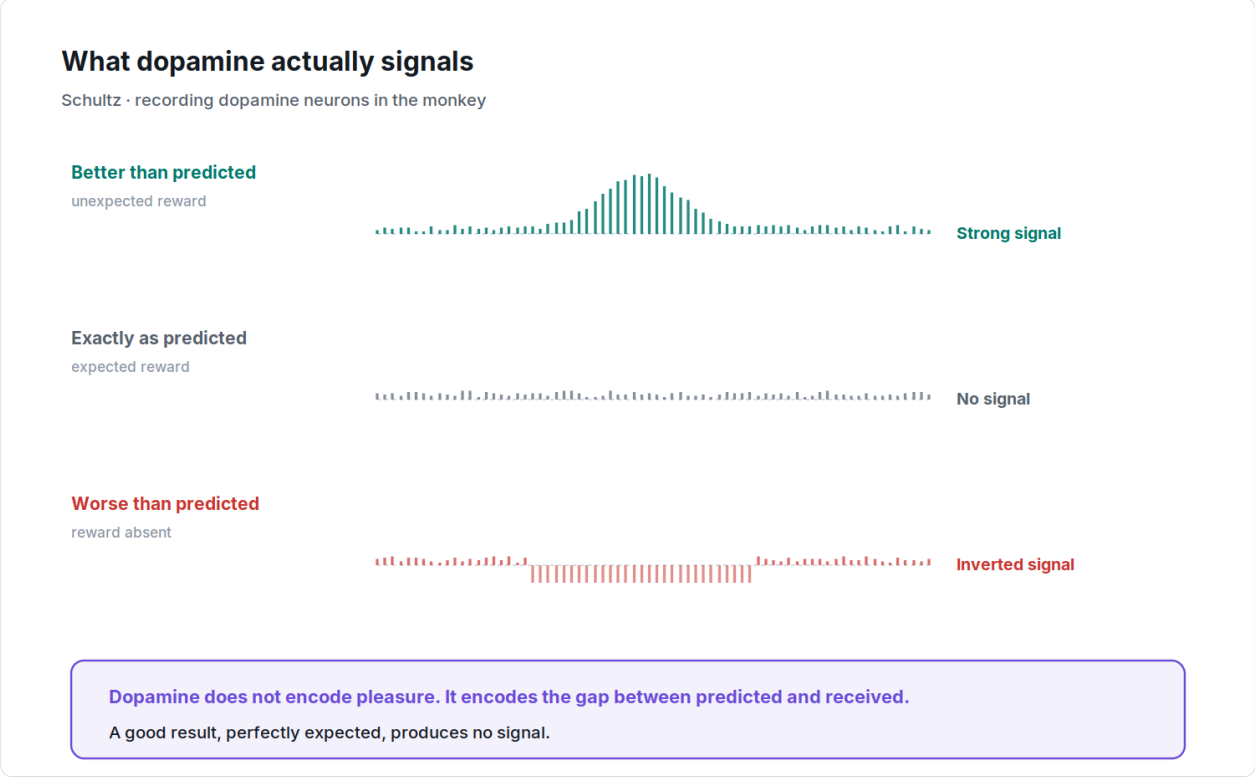
And when the juice arrives exactly as predicted? The neurons stay silent.

Schultz has described the difficulty this result posed: *why does this neuron no longer respond to the reward, when the animal is receiving it, when it responded a minute earlier?*

Answering required abandoning the pleasure hypothesis.

Three states, not one

What these neurons signal is not reward. It is the **gap between predicted and received reward**, what is called reward prediction error.



The three responses observed

Better than predicted: strong firing. This is the learning signal: something is better than you believed, update.

Exactly as predicted: baseline activity. Nothing to learn.

Worse than predicted: activity depressed below baseline. The prediction was too optimistic, correct it downward.

This mechanism is now among the best established in neurobiology. It holds in humans, monkeys and rodents.

And it has a consequence that will occupy us: **a good result, perfectly expected, produces no signal**.

Kaizen, mechanically

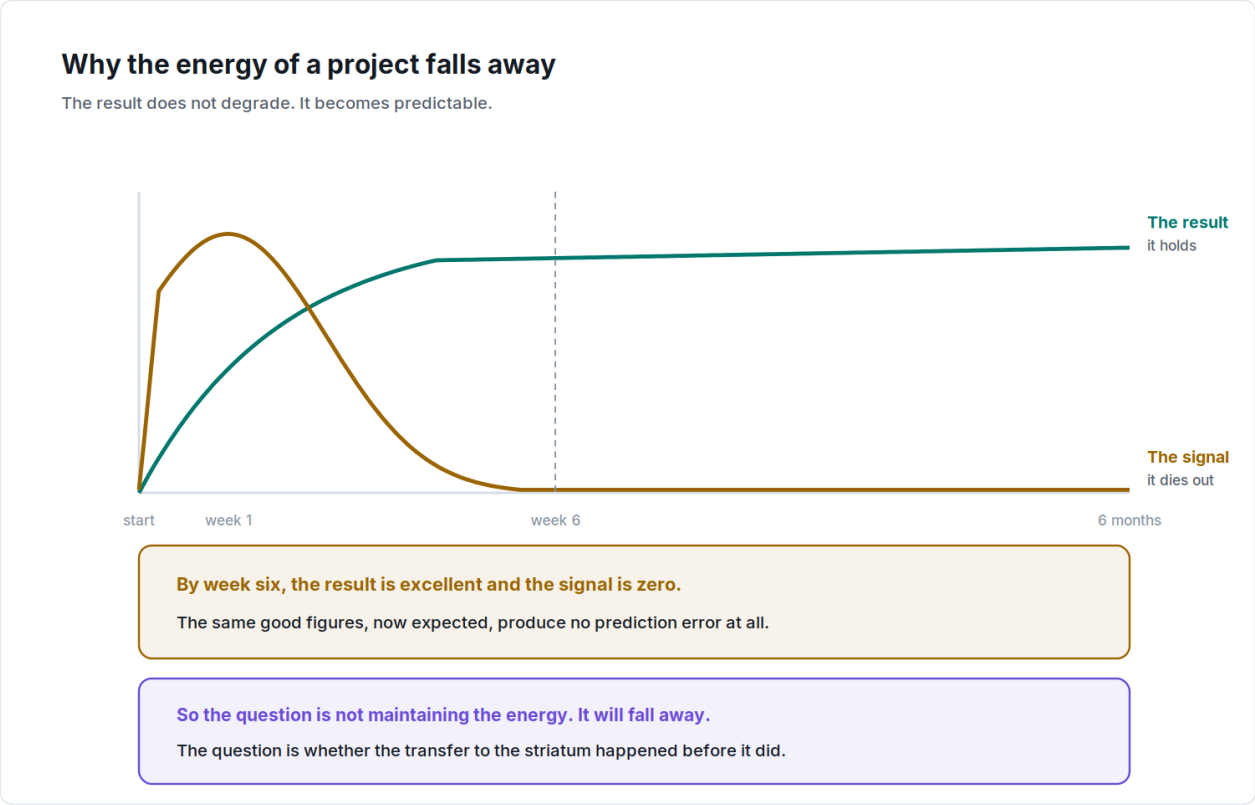
Let's revisit the improvement-project scene through this lens.

Week one. Everything is new. An idea is proposed and tested within the hour: unexpected. A bin is moved and cycle time drops, measured in front of the team: unexpected. The supervisor comes by and thanks people: unexpected.

Every event produces a positive prediction error. The energy you see on the floor isn't enthusiasm for Lean: it's a system receiving learning signals continuously.

Week six. The station runs well. Times are held. The supervisor still comes by.

But all of it is now **predicted**.



What becomes of the result and the signal

The result hasn't dropped; it's better than at the start. What has disappeared is the gap. And without a gap, no signal.

The team hasn't become less motivated. Its learning system simply has nothing left to learn from this situation.

What this changes about steering

This reading invalidates a widespread intuition: that you should "maintain the energy".

The energy is going to fall. That's mechanical, and it isn't a failure. A system that kept getting excited about a perfectly predictable result would be a malfunctioning system.

The real question is this: **did the transfer to the striatum happen before the signal died out?**

If it did, the practice continues without energy, because it no longer costs anything. That is exactly what you were after.

If it didn't, there is neither signal nor automatism left. Nothing carries the practice, and it fades.

That is why the duration of chapter 2 and the reinforcement of this chapter are two faces of one problem. You have a window, a few weeks during which the signal is still active, to install an automatism that will no longer need it.

The project isn't there to convince. It's there to hold the practice long enough for it to switch.

What this says about delayed feedback

One last point, which illuminates chapter 3.

We said that a standard whose benefit appears in a monthly indicator reinforces nothing. The mechanism now lets us say precisely why.

For an action to be reinforced, the system must be able to **attribute** the prediction error to that action. That requires temporal proximity: the signal arrives, and you know what it relates to.

A benefit aggregated over four weeks, presented in a meeting to people who weren't at the station, cannot be attributed to any particular action. It can be understood intellectually, the prefrontal cortex handles that perfectly well, but it reinforces no sequence.

Hence the recommendation of chapter 3: when a standard has no natural feedback, you have to build one. A gauge that says yes or no within a second produces an attributable prediction error, on every part.

A limit worth stating

Three cautions, and they matter, because this field is the most distorted by popular accounts.

On dopamine in general. It doesn't only encode prediction error. It is involved in motor control, motivation and learning, and its roles vary across circuits. "Dopamine, the pleasure molecule" is false, but "dopamine encodes prediction error" is a simplification of a richer system.

On the measurements. Schultz's recordings are of single neurons in monkeys, with primary rewards: juice, water. Transposing to an operator's satisfaction at an improved cycle time is a considerable extrapolation.

On managerial use. This whole chapter could be used to justify artificial gratification: surprises, variable rewards, gamification. That would be a misreading. These are the mechanisms slot machines exploit, and adults detect very quickly what is instrumental.

What remains usable without drift: **informational** and immediate feedback, the part is good or it isn't, naturally produces the prediction error that reinforces. Evaluative, delayed feedback does not.

What to take away

Dopamine **does not encode pleasure.** It encodes the gap between predicted and received.

Three states. Better than predicted: firing. As predicted: nothing. Worse than predicted: depressed activity.

The signal migrates to the cue. Once the reward is predictable, the response moves back to the first cue that predicts it.

The energy of an improvement project dies out mechanically. Not because the result degrades, but because it becomes expected.

The window is the real subject. The transfer to the striatum must have happened before the signal died out.

Feedback informs, it doesn't evaluate. Artificial gratification produces the opposite effect.

Sources

Schultz, W. (2016). *Dopamine reward prediction error coding*. Dialogues in Clinical Neuroscience, 18(1), 23–32.

Hollerman, J. R. & Schultz, W. (1998). *Dopamine neurons report an error in the temporal prediction of reward during learning*. Nature Neuroscience.

Schultz, W., Dayan, P., & Montague, P. R. (1997). *A neural substrate of prediction and reward*. Science. Formalisation of the temporal difference learning model.

On localisation: the neurons concerned lie mainly in the substantia nigra and the ventral tegmental area.

On the warning against the "pleasure molecule" reading: Schultz himself stresses that the signal is a teaching signal, not one of satisfaction.

CHAPTER 09

What circulates and isn't true

The chapter that should make you wary of this book

8 min read

In 2012, a team surveyed 242 British and Dutch teachers, teachers interested in the neuroscience of learning, so presumably well informed.

The questionnaire mixed true and false statements about the brain.

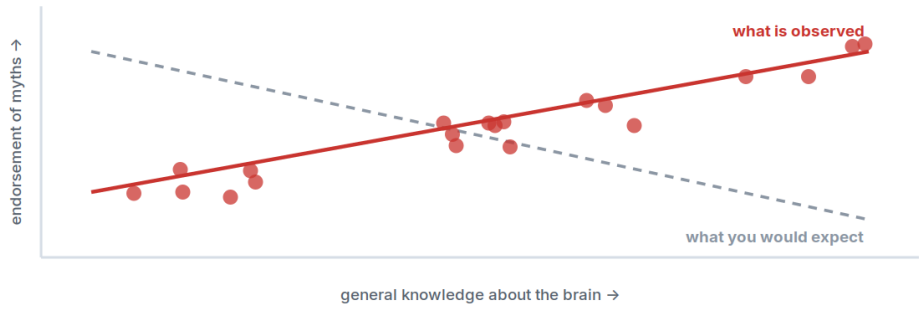
They endorsed, on average, **49% of the false statements**, while answering about 70% of general knowledge questions correctly.

That much was expected. The second result far less so.

The more you know, the more you believe

The result that should worry you

Dekker et al., 2012 · 242 teachers interested in neuroscience



The more they knew about the brain, the more they endorsed false claims.

They endorsed on average 49% of the false statements, while answering 70% of knowledge questions correctly. Interest does not immunise: it exposes.

What the study measured

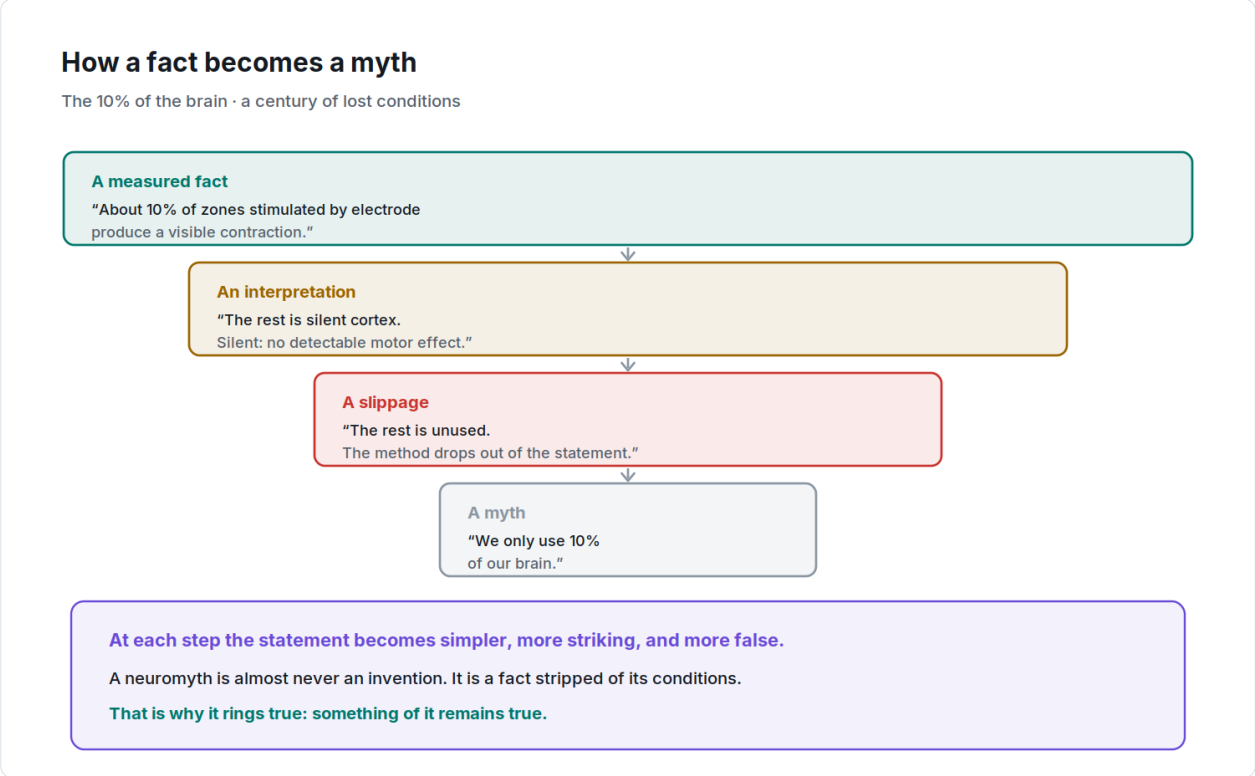
Teachers who read popular science magazines scored better on general knowledge. And **better general knowledge predicted stronger endorsement of the myths**.

The authors offer a simple explanation: people enthusiastic about applying neuroscience struggle to tell pseudoscience from fact. Interest in the field doesn't protect. It exposes.

This result applies to you, reading this book. And to me, writing it.

How a fact becomes a myth

A neuromyth is almost never an invention. It's a real fact from which, step by step, the conditions that made it true have been stripped away.



How the 10% myth was manufactured

At the start, a genuine observation. When neurosurgeons began stimulating the cortex with electrodes, over a century ago, about 10% of the zones produced a visible muscular contraction. The rest produced nothing observable **with that method**.

They spoke of "silent cortex". The term was honest.

Then silent became unused, the method dropped out of the statement, and what remained was a striking phrase still sold today.

This mechanism applies to almost every statement that follows.

Six statements, and what they're worth

"Everyone has their learning style"

The most widespread of all, endorsed by a large majority of the general public.

What's true: people have preferences, and express them fairly stably.

What isn't: that teaching to those preferences improves learning. Studies testing precisely that hypothesis find no such advantage.

What does work is **multiple representation**: presenting the same thing in several ways, for everyone. That isn't the same recommendation.

"Left brain, right brain"

What's true: the hemispheres show specialisations. Language draws more on the left in most right-handers.

What isn't: that there are "left-dominant" analytical individuals and "right-dominant" creative ones, and that this typology predicts a way of learning.

It is the foundation of many profiling tools sold to companies.

"21 days to form a habit"

Covered in chapter 2. Origin: a 1960s plastic surgery book. Measured data: 66 days at the median, from 18 to 254.

This is probably the costliest myth on the list for a practitioner, because it is what justifies three-week support.

"Dopamine is the pleasure molecule"

Covered in chapter 8. It encodes the gap between predicted and received, not satisfaction. Schultz himself insists on the point: it is a teaching signal.

The distorted version leads to gratification schemes that miss their target.

"Stress hijacks the brain"

The popular version of chapter 6's mechanism, and it is too clean.

What's documented: a relative weakening of prefrontal control and a relative strengthening of subcortical responses.

What's overstated: talking about "shutting down" the prefrontal cortex, or an "amygdala hijack", as if it were a switch.

"The brain rewires itself constantly"

What's true: plasticity exists throughout life.

What's distorted: the pace and the voluntarism. What is observed is slow, gradual, largely involuntary, and doesn't reverse on command.

This version feeds the idea that wanting is enough to change, the exact opposite of what this book argues.

Why it matters

These myths sell things. The most widespread are precisely those tied to commercial programmes: profiling tools, "brain-based" training, hemispheric integration methods.

They steer decisions. Sorting operators by profile, planning three weeks of support because twenty-one days would do, are choices with consequences.

They discredit the rest. This is the worst of it. When an engineer hears "right-brain profile" in a training session, they rightly file the whole discourse under nonsense, including the parts that stood up.

Three questions that filter

On whom, and measured how? A statement mentioning neither method, nor sample, nor conditions has probably lost what made it true.

How far from the result? There's a distance between "activity in this region increases during this task" and "this is why your team resists". The second may be defensible, but it crosses several levels, and whoever says it should say so.

Does it sell something? Not disqualifying in itself. But a statement that conveniently leads to a proprietary tool deserves an extra check.

> *The right question isn't "is it true?" but "what was measured, on whom, and what remains when you put the conditions back?"*

What this says about this book

Let me apply those three questions to my own text.

On the first. Every claim is referenced with author, year and journal. But I don't have access to full paywalled texts: I work from abstracts and syntheses. I may have missed a methodological restriction.

On the second. That is the book's structural weakness, and the next chapter is devoted to it. The whole exercise consists of crossing the distance between a laboratory result and a shop floor.

On the third. This book is written by someone whose profession is continuous improvement, and it lends credibility to an approach I practise. That's an interest, and you should factor it in.

What to take away

Interest in the subject doesn't immunise you. Knowing more about the brain predicted stronger endorsement of the myths.

A myth is born of a fact stripped of its conditions. That's why it rings true.

Learning styles are the most widespread myth. Preferences exist; the advantage of teaching to them isn't established.

Two of this book's mechanisms have distorted versions. Dopamine "of pleasure" and the brain "hijacked" by stress.

Three questions filter. Measured how? How far from the result? Does it sell something?

Sources

Dekker, S., Lee, N. C., Howard-Jones, P., & Jolles, J. (2012). *Neuromyths in education: Prevalence and predictors of misconceptions among teachers*. Frontiers in Psychology, 3, 429.

Macdonald, K. et al. (2017). *Dispelling the myth: Training in education or neuroscience decreases but does not eliminate beliefs in neuromyths*. Frontiers in Psychology, 8, 1314.

Howard-Jones, P. A. (2014). *Neuroscience and education: Myths and messages*. Nature Reviews Neuroscience, 15(12), 817–824.

OECD, *Brain and Learning* project (2002). First institutional warning about neuromyths.

On the origin of the 10% myth: early twentieth-century cortical stimulation work and the notion of "silent cortex".

CHAPTER 10

What this book does not prove

Where the weaknesses are

6 min read

Every mechanism in this book was established in a laboratory.

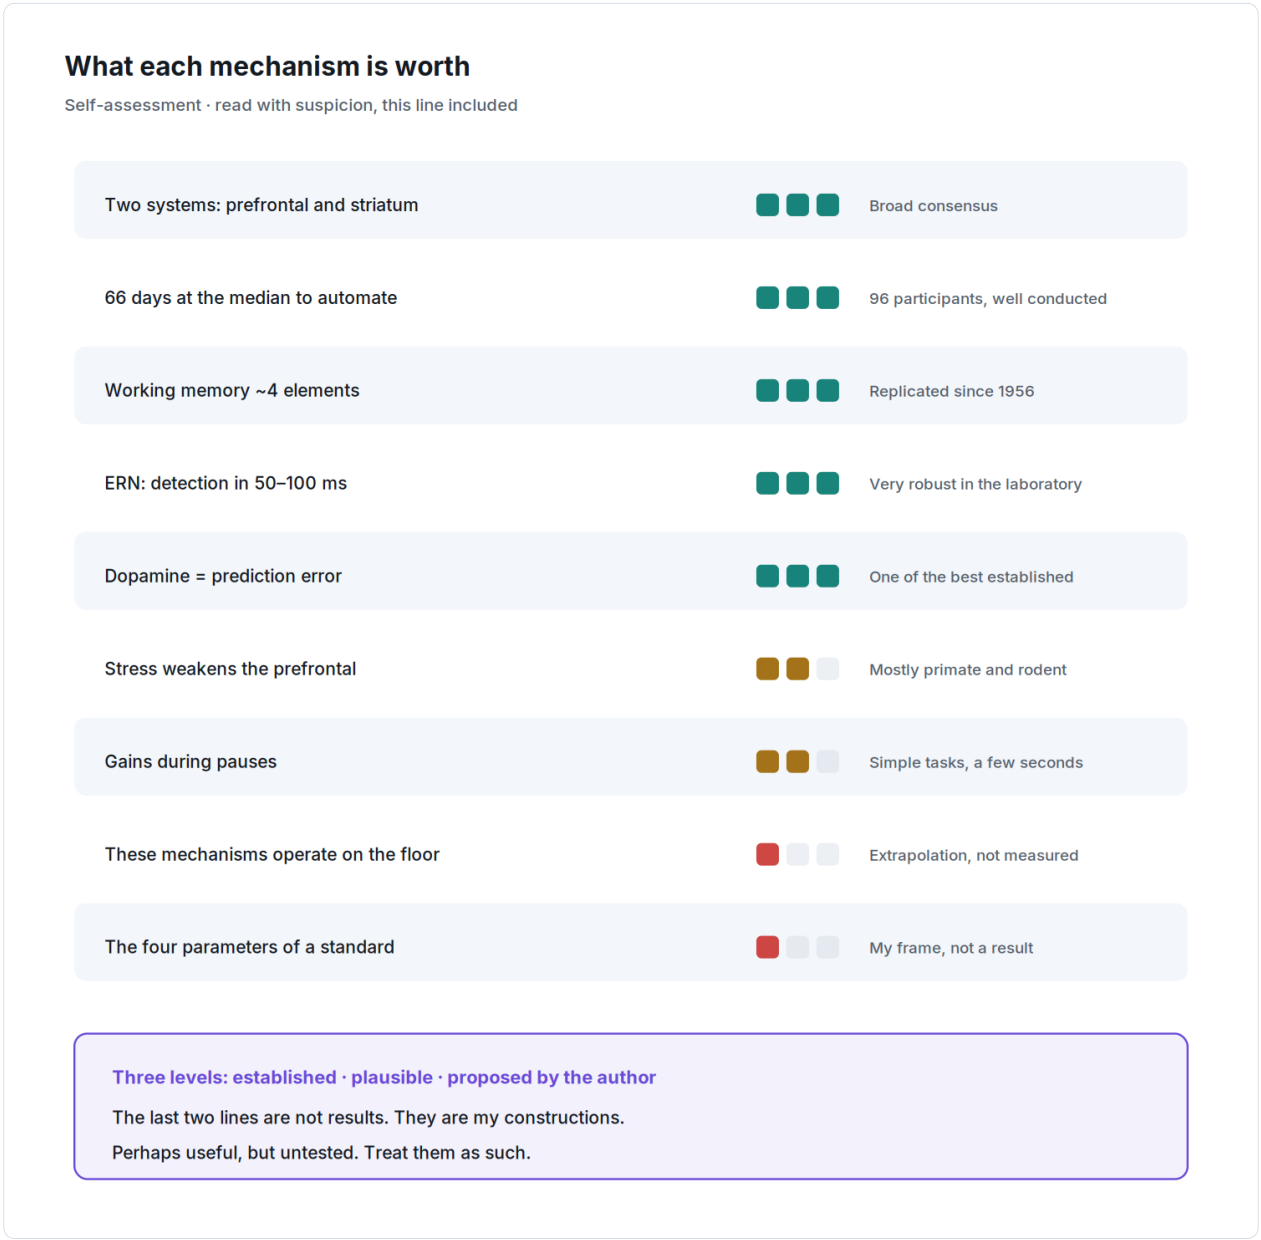
On tasks lasting seconds. With one variable isolated. Often in animals. Always in a controlled environment.

A shop floor is the exact opposite: everything varies at once, nothing is isolated, and nobody measures anything to the millisecond.

Every chapter of this book crosses that distance. That is its purpose, and its weakness.

Mechanism by mechanism

Rather than leave this vague, here is an explicit assessment.



What each mechanism is worth

Three levels, defined precisely.

Established: reproduced result, consensus in its field, and I transpose it without changing its nature.

Plausible: the original result is solid, but I apply it to a different context, longer timescales, social stakes, an uncontrolled environment.

Proposed: this is my construction. The four parameters of chapter 3 are not a research result. They are a frame I assembled from elements drawn from separate studies. It may be useful; it has not been tested.

That last category deserves attention, because it passes easily for science. A four-point frame with references at the foot **looks** like a result.

The leap I make everywhere

Take the clearest case.

The ERN of chapter 5 is measured on tasks of a few hundred milliseconds, with consequence-free errors, in students sitting at a screen.

From that I conclude that the operator fitting a part backwards detects the error before becoming aware of it.

That's plausible: error detection is a general action-monitoring mechanism. It is not demonstrated in that context.

The same applies to dopamine, measured on primary rewards in monkeys, which I transpose to an operator's satisfaction at an improved cycle time.

Three further limits

I don't have access to full paywalled texts. I work from abstracts and syntheses. That's enough to render a main result faithfully, not always to spot a methodological restriction.

No company data. This book contains no measurement taken in a real organisation. The examples are constructed illustrations, plausible but proving nothing.

No economic or social dimension. A transformation unfolds in a context of power relations, sometimes of job cuts. This book doesn't address that, which makes the mechanisms described look more decisive than they are.

How to read this book

Not as a demonstration, but as a **shift of attention**.

Its value doesn't rest on the solidity of each isolated claim. It rests on offering a different reading of a familiar failure: instead of "they didn't want to", look at the signal, the load, the duration, the feedback.

That shift stays useful even where the mechanisms are imperfectly established, because it leads you to examine things that are **modifiable**.

And it is falsifiable, which is its best guarantee. If you install a signal at the workstation, support for ten weeks instead of three, cut the standard to three points, build immediate feedback, and nothing changes, then this book is wrong for your situation.

That is more honest than a promise you cannot test.

What to take away

Distance is the main weakness. From laboratory to shop floor, each step removes conditions.

Some claims are mine. The four parameters of chapter 3 are a proposed frame, not a result.

The best-established mechanisms are the most general. Two systems, working memory, prediction error. The transpositions are more fragile than the facts.

The book holds as a shift, not as a demonstration. And it is falsifiable.

Sources

Open Science Collaboration (2015). *Estimating the reproducibility of psychological science*. Science, 349(6251). On the uneven reliability of the psychological literature.

The sources for each mechanism appear at the end of their respective chapters.

CHAPTER 11

What changes on Monday

Seven questions, none about people

5 min read

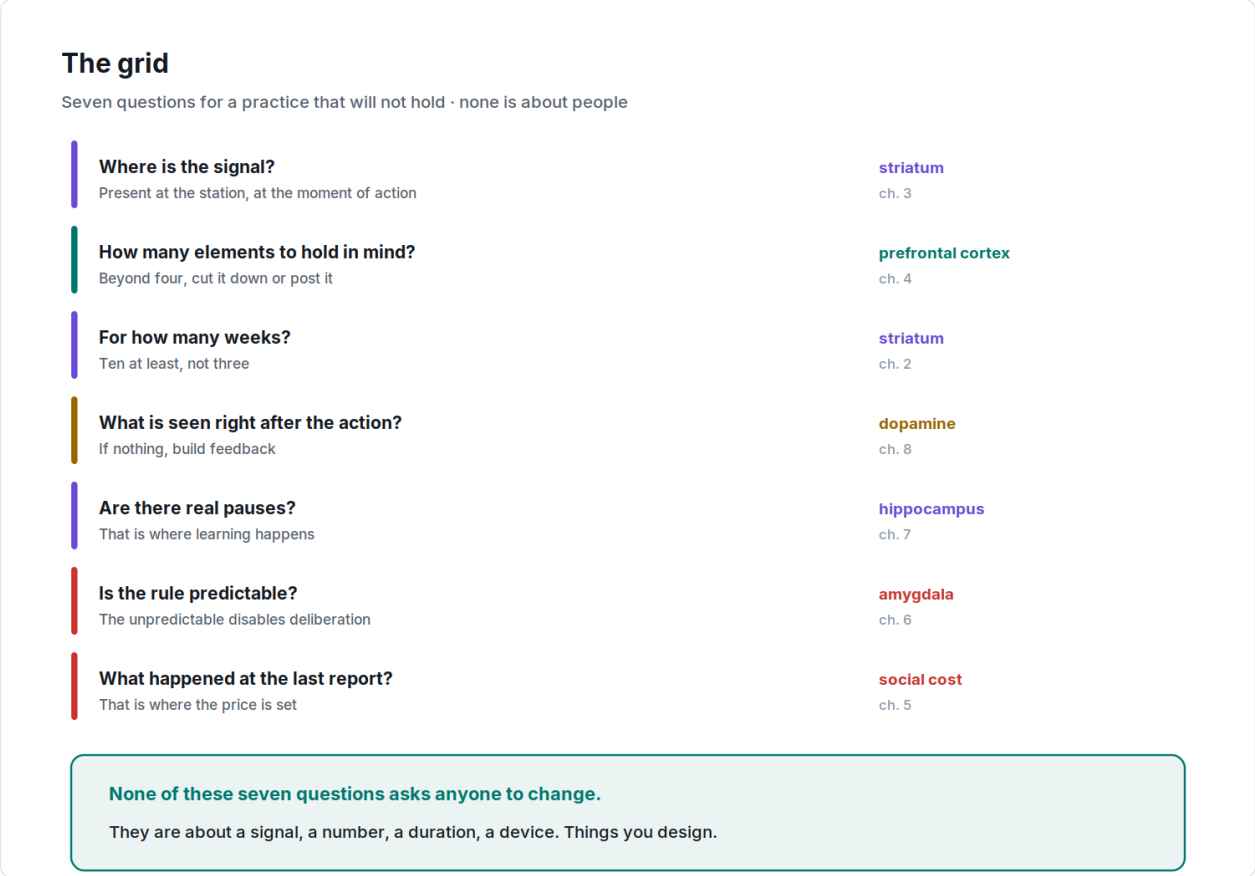
Back to the scene from the introduction.

You walk through a workshop six months after an improvement project. The markings are worn, the board shows seven-week-old figures, the station runs as it did before.

This book hasn't claimed to give you a method for avoiding that. It has offered a different way of looking at the same scene.

Instead of asking who failed to hold, ask **which mechanism wasn't respected**.

Seven questions



The seven questions and their mechanisms

Take a practice that isn't settling with you right now. Run it through these seven questions.

In most cases, two or three answers will be missing, and they will point precisely at what to change.

Note what none of them asks: none is about people. None looks for who is on board and who is dragging.

They are about a signal, a number of elements, a duration, a device, a rule. Things you design.

That is what makes them usable. A recommendation requiring people to become different isn't a recommendation, it's a wish.

Five substitutions

Nothing that follows needs a budget or an approval.

The standard in the meeting room becomes the standard at the workstation. The striatum receives a signal instead of waiting to be thought of.

Three dense days of training become short sessions separated by nights. The hippocampus can consolidate.

The monthly indicator becomes immediate feedback. Dopamine can attribute the gap to the action.

Three weeks of support become ten. The transfer to the striatum has time to happen.

A rule that varies becomes a rule stated in advance. The situation becomes controllable again, and the prefrontal cortex stays available.

The levers that matter are free. They belong to design, not to conviction.

Three reflexes that change

You no longer launch a project without knowing who will hold the follow-up. A five-day project with no support over ten weeks is an investment whose outcome you already know. Better to launch two and hold them than six and abandon them.

You no longer read silence as an absence of problems. Detection is automatic in everyone. What varies is what it costs to speak. The only useful indicator: what happened the last time someone reported something?

You no longer ask people why they do what they do. The question puts them on the defensive, and a person on the defensive describes what protects them. What stays useful is direct observation.

What doesn't change

The tools remain necessary. Nothing here suggests abandoning SMED or VSM. They are good; they aren't sufficient.

The floor remains where truth lives. None of these mechanisms shows up on a dashboard.

People are not systems. This book describes mechanisms because that is what design can modify. But a team isn't reducible to them, and a manager treating colleagues as circuits to optimise would miss what matters.

The shift proposed here runs from "it's their fault" to "it was badly designed". It doesn't run from "these are people" to "these are circuits".

One last thing

> *A culture of continuous improvement cannot be decreed and does not depend on people's conviction. It settles slowly, under precise conditions, and disappears as soon as those conditions stop being met.*

That sentence is less inspiring than most you'll find in a conclusion. It has one advantage: it points at things you can do.

The signal, the load, the duration, the feedback, the pauses, the predictability, the cost of speaking up.

Seven parameters, all modifiable, nearly all free.

It isn't much. It is also, probably, what separates the initiatives that hold from those that fade.