TULSER

No. 03 · © 2026 Tulser

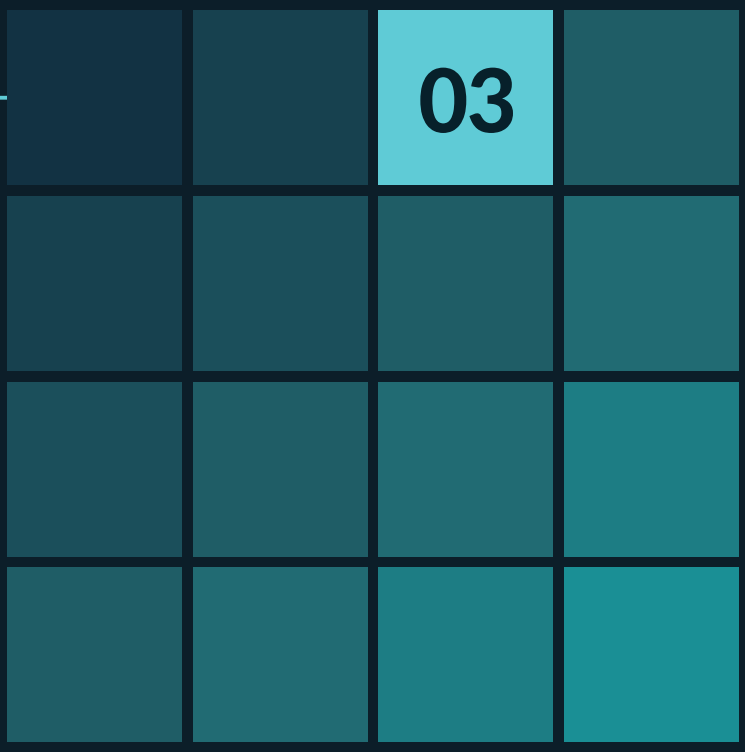
WHITE PAPERS · PERFORMANCE CONSULTING

BUSINESS PROBLEM

*A result is moving the wrong way.*

# Environment before *capability*

*What the largest synthesis of behavioral evidence says about where performance comes from, and why expertise is the last resource still rationed.*




JOS ARETS

*with Vivian Heijnen*

IN THIS PAPER

# Telling people more *changes almost nothing.*

| | |
|---|---|
| THE QUESTION | *A number is moving the wrong way and the obvious fix is to inform, persuade or train. What does the evidence say that actually changes behavior at work?* |
| FOR | Line managers and directors, and the staff functions asked to fix performance: OD, quality, continuous improvement, HSE, HR business partners. |
| READ | 16 minutes · with a worked case and a one-page instrument. |


| | | |
|---|---|---|
| 01 | The reflex, and why it is reasonable | Knowledge and attitude are visible, and a program fits them. |
| 02 | What the evidence actually says | Nine intervention targets, ranked, and the two distinctions that keep it honest. |
| 03 | Five resources, one word | What “access” means once you look at a working day, and how to test each one in a day. |
| 04 | Two engineers and three regions | A worked case in which the missing resource is not knowledge. |
| 05 | The last scarce resource | Expertise has always been rationed by place and seniority. |
| 06 | What this does not mean | Five limits, because the rest is not credible without them. |
| · | Run it yourself | Which of the five is missing, on the problem in front of you. |

> *Policymakers should focus on interventions that enable individuals to circumvent obstacles to enacting desirable behaviours rather than targeting salient but ineffective determinants of behaviour such as knowledge and beliefs.*

**ALBARRACÍN, FAYAZ-FARKHAD & GRANADOS SAMAYOA**

*Nature Reviews Psychology, 2024*

# The reflex,
## *and why it is reasonable*

When a result moves the wrong way, the first instinct is to inform, to persuade, or to train. That instinct is not lazy. It is the one option that is visible, purchasable and available this week, and every organization is built to supply it.

Look at what the alternative would require. Changing what someone can reach at the moment of the decision means touching a process, a system, an authorization, a rota or an incentive: each owned by somebody else, each with its own committee. Changing what someone knows requires a supplier, a date and a room. One of those can be arranged this quarter. The other cannot.

So the request writes itself, and it is sincere. The manager has watched the problem for six weeks, has a theory, and puts the theory in the only form the organization reliably accepts. Nobody in this chain is being careless. They are choosing the intervention that the system makes available.

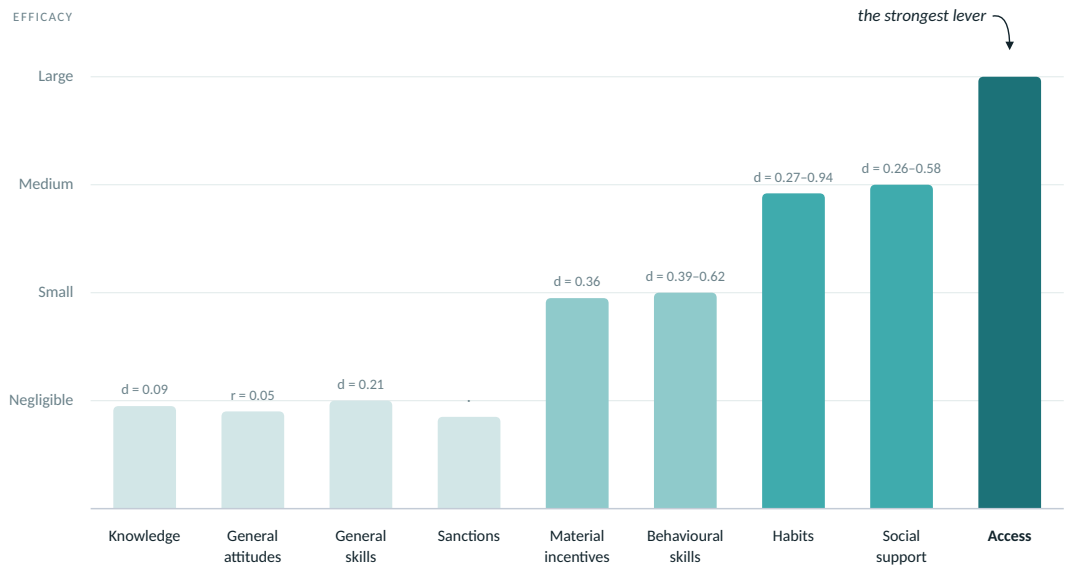
> **THE DISTINCTION THAT DECIDES THE MONEY**
>
> Knowledge is something you give a person. Access is something you arrange around them. The first is a transaction and can be scheduled; the second is a change to the work and has to be negotiated. Which is precisely why organizations buy the first and hope it substitutes for the second.

The reflex has a long record of being named and surviving anyway. Harless told the profession in 1985 what selling training without a front-end analysis amounted to, in the society's own journal and in language nobody could mistake: number 02 quotes the sentence. Four decades later the request still arrives with the remedy attached.

The uncomfortable part is not that this is wrong in principle. It is that the size of the effect is now measured, across hundreds of studies and dozens of domains, and it is very small.

# What the evidence *actually says*

In 2024 Dolores Albarracín and colleagues did something the field had not done before. Rather than run another study, they gathered the available meta-analyses of behavior prediction and intervention efficacy across domains, and ranked what actually moves behavior and what does not.




Nine targets an intervention can aim at, ordered by how much behavior actually changes when you aim at them: the four on the left are what organizations buy most. From Albarracín, Fayaz-Farkhad and Granados Samayoa (2024), synthesising the available meta-analyses across domains; bar height shows the reported band and darkness carries the same information.

Read that figure from left to right and it inverts the way most performance budgets are built. Handing people information has a negligible effect. Changing what people think in general has a negligible effect. Teaching a broad, non-specific skill has a negligible effect. So do rules and sanctions, which is the other reflex, if information will not do it, forbid the behavior.

Effects only start to appear once an intervention stops addressing the mind and starts addressing the situation. Money helps a little. A concrete, behavior-specific skill helps a little more. A changed routine helps more again. Arranged help from colleagues helps more still. And the largest effects belong to interventions that remove whatever was standing between the intention and the act.

# Two distinctions *that keep it honest*

A figure this convenient invites two objections, and both are worth answering before anyone builds a decision on it.

### A DETERMINANT IS NOT AN INTERVENTION TARGET

The strongest single *predictor* of how well someone performs is not in that figure at all. In the training-transfer research it is cognitive ability, which tracks how much of a program survives contact with the job more closely than any feature of the workplace does (Blume et al., 2010). That does not weaken the argument. You cannot buy an intervention that raises cognitive ability. Something that explains why people differ, but cannot be moved, is useful for selecting people and useless for improving performance.

### THE SAME RESULT, FOUND INSIDE ORGANIZATIONS

Albarracín's evidence comes from health and policy behavior. The parallel research inside organizations reaches the same place, and it has been reaching it for fifteen years. Pooling eighty-nine studies of training transfer: whether the manager backs it, and what the climate is like in the place people return to, both track how much of a program survives, and the work environment weighs six times more heavily on open-ended work than on tightly scripted work, because the more discretion the job allows, the more the surroundings decide what happens (Blume et al., 2010). Support from peers, manager and organization together accounts for roughly a third of the difference in whether training transfers at all (Hughes et al., 2020), which is to say that the environment decides most of what a program achieves, and it keeps deciding long after the room has emptied.

#### THE FIGURES, FOR ANYONE WHO WANTS THEM

Blume and colleagues, a meta-analytic review of transfer of training: cognitive ability correlates with transfer at $r = .37$, supervisor support at .31 and transfer climate at .27, across 89 studies. The work-environment effect is roughly six times larger for open skills than for closed ones. Hughes and colleagues: peer, supervisor and organizational support together account for about a third of the variance in transfer sustainment.

#### WHAT THAT MEANS FOR A TRAINING BUDGET

Two organizations can run the identical program, with the identical designer and the identical participants, and get different results, because transfer is decided after the room empties, by conditions nobody costed. Which is why "was it a good course?" is close to unanswerable, and "what did the work allow afterwards?" is answerable in a week.

# And none of it *is new*

None of this is new. Gilbert set the priority in 1978: “only when we have made a proper analysis of accomplishments and their measures will we have any sensible reason to concern ourselves with behavior” (Gilbert, 1978, p. 73). Rummler established where those accomplishments leak: the seams between departments (Rummler & Brache, 2013). Three disciplines, three routes, one conclusion.

*The knowledge was never the scarce thing.*
*Access was.*

# Five resources, *one word*

"Access" sounds abstract until you watch a working day. Then it resolves into five concrete provisions, each of which someone can arrange, and each of which is quietly missing somewhere in every organization. The five are our own grouping of what analyses keep turning up, not a published taxonomy; what they are worth is in the middle column.

| RESOURCE | THE QUESTION | WHAT ITS ABSENCE LOOKS LIKE |
|---|---|---|
| **Information** | Is what I need to decide in front of me when I decide? | Data that arrives a day late, or lives in a system I cannot open. |
| **Tools** | Does the equipment support the work or fight it? | Workarounds, second spreadsheets, steps done twice. |
| **Time** | Is there room to do it properly? | The step everyone agrees is necessary is the first one dropped when the day runs late. |
| **Authority** | May I act, or must I ask? | Decisions waiting for someone two levels up who is traveling. |
| **Expertise** | Can I reach someone who has seen this before? | Long resolution times far from head office; the same problem solved slowly, repeatedly. |

Five resources. The first four are usually arranged locally. The fifth almost never is. *[our formulation, drawn from client work; not a published taxonomy]*

Four of the five behave alike. They are visible once you look, they belong to someone who can be named, and they can be changed within a quarter by a manager with a modest mandate. Unclear information, missing tools, absent time and withheld authority are the ordinary furniture of performance analysis, and the fixes for them are ordinary too.

The fifth behaves differently, and it is the one this paper is really about. Expertise cannot be bought in a quarter, cannot be moved to where it is needed, and does not scale by hiring, because the people who have it are the same people whose time you are trying to protect. Every organization rations it, and almost none of them decided to.

That rationing is not a policy. It is geography and seniority. Whoever sits near the person who has seen the problem before gets an answer in ten minutes. Whoever sits elsewhere gets a manual. The difference between those two people is not competence, and it is not attitude. It is what they can reach.

# How to test each one
## *without running a project*

A list of five is a checklist, which is worth little. What makes it an instrument is knowing how to test each one, what a positive test looks like, and roughly what it costs to fix. None of these five tests takes more than a day.

| RESOURCE | THE TEST | ROUGHLY WHAT A FIX COSTS |
|---|---|---|
| **Information** | Sit with somebody at the moment they decide and ask what they would want to know. Then find out how long it takes them to get it, in minutes. | Usually small. A report on a schedule, a field on a screen, a standing item. Days of work, not months. |
| **Tools** | Ask what they built themselves. Every private spreadsheet, laminated card and saved search is a tool the organization did not supply. | Small to moderate. Often the tool exists and has not been distributed: a license rather than a purchase. |
| **Time** | Take the step everyone agrees is necessary and ask when it is dropped. If the answer is “when the day runs late”, time is the binding condition. | Large and political. Time comes out of something else, and the something else has an owner. |
| **Authority** | Follow one decision and count how many people have to touch it. Then ask what the largest thing is that the person doing the work may decide alone. | Small in money, large in nerve. It costs nothing to raise a threshold and somebody has to be willing to. |
| **Expertise** | Ask who they ring when they are stuck, how long it takes, and what they do when that person is unavailable. The last answer is the finding. | Historically very large, which is why it is rationed rather than solved. |

Five tests, a day each. In our practice the third and fifth turn out to be binding most often, and they are the two that cost most to fix.

The cheap fixes are at the top, which is the opposite of where attention usually goes: an organization will run a development program before it will put a field on a screen.

HOW YOU KNOW YOU WERE RIGHT

Each of these has the same falsification test. Arrange the resource for one team and not for a comparable one, wait a month, and compare. If the measure moves only where it was arranged, the condition was binding. If neither moves, it was not: the cheapest wrong answer there is.

# Which of the five *is actually binding*

Usually more than one is missing and only one is binding: the one that, if you arranged it tomorrow, would move the measure on its own. Finding which is a matter of sequence rather than judgement.

| THE MOVE | WHAT IT ESTABLISHES |
|---|---|
| **1 · Ask the person who already achieves it** | Which of the five they have that the others do not. This is number 07's subtraction, and it names the candidate in an afternoon. |
| **2 · Ask the person who does not** | What stopped them, on the last occasion, in that order of words. People remember obstacles precisely and remember routine badly. |
| **3 · Check whether the candidate varies with the measure** | If the teams with the resource perform better and the teams without it do not, the candidate is doing work. If both look the same, it is present but not binding. |
| **4 · Arrange it for one group** | The cheapest of the candidates, for one comparable team, for a month. Then compare. |

Four moves, roughly two weeks. Step three eliminates most candidates at no cost, which is why it comes before step four.

Step two is the one people skip, and it is free. The exemplary performer has arranged the resource and stopped noticing it; the standard performer is standing in front of its absence every day and can describe it exactly. Ask one what they have and the other what stopped them, and the two answers name the condition between them.

### WHEN MORE THAN ONE IS BINDING AT ONCE

It happens, and it is why interventions that address a real condition sometimes produce nothing. If information and authority are both missing, giving somebody the number they need does not help if they still cannot act on it. The test in step four catches this: arrange one, see nothing, and the conclusion is not that the condition was wrong but that it was not sufficient. Add the second and test again before concluding either was a mistake.

04 · A WORKED CASE

# Two engineers
## *and three regions*

A manufacturer runs the same production line at its main plant and at three regional sites. Same equipment, same procedures, same certification for every technician. And the same fault takes four hours to clear at the main plant and two and a half days everywhere else.

The reading offered by the situation is the familiar one. Regional technicians must be less experienced, or less thorough, or less committed. The request that follows is a three-day technical refresher for sixty people, and it is entirely defensible: nobody has a better explanation, and training is the explanation the organization knows how to buy.

| WHAT IT COSTS | VOLUME | RATE | PER YEAR |
|---|---|---|---|
| **Extra downtime at regional sites** | 4 events per month × 20 hours. | $1,500 lost contribution per hour. | **$1.44M** |
| **The requested refresher** | 60 people × 3 days. | $50 loaded hour · $400 direct per day. | $144k, once |

Worked scenario. Every figure recomputes from the volumes and rates shown; the proportions carry the argument, not the currency.

Now the diagnosis. The technicians are certified and the procedures are identical, so knowledge is unlikely to be the constraint, and a two-minute check confirms it: on the rare occasions a regional team reaches the main plant by phone, the fault clears in about the same four hours. What differs is not what people know. It is that two engineers who have seen this failure two hundred times sit at the main plant, and everyone else has a manual that describes the procedure but not the reasoning.

The missing resource is access to expertise. And note what the refresher would have produced: sixty well-trained technicians returning to sites where the same two engineers are still unreachable, and a manager concluding a year later that the training did not stick.

The gap was never in what they knew.
*It was in who they could reach.*

# The last *scarce resource*

Information became inexpensive, then abundant, then free. Tools became rentable by the month. Authority can be delegated by writing it down. Expertise allowed none of this, and until very recently there was no reason to expect it to.

Expertise is different because it is not a document. What the two engineers have is not the manual; it is the reasoning that decides which of nine plausible causes to test first, built from two hundred encounters with the same machine. That reasoning has always traveled in exactly two ways: by standing next to the person who has it, or by spending years acquiring it yourself. Both are expensive, both are slow, and both are strictly limited by where you happen to sit.

Which is why every organization, without deciding to, ends up with a competence gradient that follows its floor plan. The head office is where the answers are. The distance from head office is, in practice, the resolution time.

### AND THE RESOURCE THAT HAS JUST STOPPED BEING SCARCE

That sentence, *until very recently*, is doing work. Expertise is the one resource on this list whose economics changed in the last few years, and the change is not that people got better at reasoning. It is that reasoning built over two hundred encounters can now be written down in a form that answers questions, holds a sequence and travels to the branch that is four hours from head office.

This paper does not pursue that, because it is an argument about where performance comes from and not about tools. **Number 15 pursues it, and turns this same argument on the discipline itself:** a method that has been right for sixty years and rationed for sixty years by exactly the two channels named above. What happens to a competence gradient when the fifth resource stops being scarce is the last question this series asks.

# What this *does not mean*

An argument this convenient for one side deserves its boundaries stated first. Five of them, plainly.

- **It is not an argument against learning.** Where the cause genuinely is a shortfall in knowledge or skill, well-designed learning is exactly right and nothing substitutes for it. The evidence does not say teaching never works. It says teaching is a weak lever when the obstacle is not ignorance, which is most of the time, and is a claim about frequency, not about value.
- **Training effects cannot be isolated anyway.** Reviewing forty years of organizational-level research, Garavan and colleagues conclude that training cannot be separated from the system it sits in: the effects are interconnected, time-dependent and bidirectional (Garavan et al., 2021). That cuts both ways here. It supports the argument that the environment dominates, and it forbids anybody, including us, from settling the question with an effect size.
- **The evidence comes from another domain.** Albarracín and colleagues synthesized health, policy and social behavior, not work organizations. The ranking is stable within that body; carrying it across to a production line or a finance department is an inference rather than a finding, and this paper says so rather than dressing it as a result. It is a well-supported inference: Gilbert and Rummler arrived at the same ordering from inside organizations, by different routes and without the psychology.
- **Access without competence is noise.** Reaching an expert helps only someone who can act on what they hear. The five resources are conditions, not substitutes for one another; a technician who has never been certified is not helped by a faster answer.
- **A system that carries reasoning can be wrong.** It should therefore say how confident it is and what would raise that confidence, and the decision should stay with the person who owns the consequence. A tool that hides its uncertainty is a worse adviser than a manual.

# And what to do
## *about the one that is not settled*

The third limit is an open debate rather than a settled question, and it deserves to be treated as one.

Whether a ranking established in health, policy and social behavior carries over to work organizations has not been tested directly. It could be. A replication in an operational setting (same design, outcome measures taken from the work rather than from self-report) would settle in one study what currently rests on inference, and the field would be better for it. Until somebody runs it, the defensible position is that the direction is well supported and the size of the effect in a workplace is not.

### WHAT IT IS WISE TO TAKE ACCOUNT OF MEANWHILE

Three conclusions hold whichever way the question resolves. That examining the five resources is cheap and buying instruction is not, so looking first costs a week and can save a program. That two traditions working inside organizations (Gilbert on conditions, Rummler on levels) arrived at the same ordering independently, which is weaker evidence than a replication and stronger than none. And that an analysis which has ruled the conditions out makes a far better case for the training than one that never looked, so the week is not wasted even when the answer turns out to be instruction after all.

### SUMMARY

Nothing here says people do not matter. It says that when a result is moving the wrong way, the conditions around the work are the more likely cause and the less costly fix, and that of those conditions, the one nobody has been able to arrange is the one now becoming arrangeable.

RUN IT YOURSELF

# Which of the five *is actually missing?*

Fifteen minutes, on one problem you own. Take the group whose performance is below where it needs to be, and answer for them rather than about them.

**1 · The result and the number**

*What is moving, from what to what, and who reports it?*

**2 · Someone who already achieves it**

*Name a person, shift, site or unit getting the result under the same conditions.*

**3 · The five, one line each**

*Information, tools, time, authority, expertise, which does that person have that the others do not?*

**4 · The least costly of the five to arrange**

*Rank by what it would cost to fix and how much of the gap it explains. Start where those overlap.*

**5 · What the proposed solution would cost**

*Direct spend, travel and working hours, set against what the gap in row 1 is costing. Number 13 does that arithmetic.*

WHERE THIS PAPER SITS

# Two problems, *one method, sixteen papers*

| | BUSINESS PROBLEM | TRAINING REQUEST |
|---|---|---|
| **Owned by** | The manager or director who owns the result. | L&D and the HR business partner. |
| **What it prevents** | Spending on the wrong intervention: a system, a reorganization, a hire. | Spending on a program that cannot reach the cause. |
| **What it returns** | The fix that moves the number, chosen on evidence. | The budget and the working hours that were not spent. |
| **Method used** | The same three analyses and five questions, which is why one capability serves both, and why the diagnosis no longer has to be bought in case by case. | |

One method, two returns, and a standing capability instead of a standing consultancy line.

*Two models carry the method. Gilbert's six conditions ask what kind of cause; Rummler and Brache's three levels ask where it sits. Both are sound and this series uses both, and it leads with Gilbert, because his unit is the accomplishment, which is what an organization already counts, and because his six conditions are the categories an organization can actually change.*

# Where the method *comes from*

*Two further pieces close the set and carry no number, because each points at all the numbered papers. Before the guide sits Arets L&D Business Models, the destination piece for the people who receive the requests: four models a learning function can run, and the route from HR-driven to business-driven. Last sits Solutions that cannot work, the field guide, ordered by the sentence that lands on your desk: seventeen forms a performance problem arrives in, each with an address in one of the numbered papers. Three orders in all: the destination you choose, the question you are on, and the form that landed on your desk.*

## HOW THIS SERIES CAME ABOUT

Human performance technology travels the way a craft does: from hand to hand. Vivian Heijnen and Jos Arets learned the analysis from Dennis Mankin, in San Diego, more than twenty years ago, and have kept in contact with him ever since. Both remain grateful to him: for the rigor of the method as he taught it, and for the leadership he has given performance consulting over a long career.

What followed was twenty years of use. Vivian took the method into organizations in the Netherlands and well beyond them, and what accumulates over that many analyses is largely tacit: the cut nobody thought to make, the question that opens a defensive room, the finding that turns out to be an artifact of the data. Making that knowledge explicit, case by case, is what this series is for. Jos wrote it down as it went, in two books and more than forty articles and blog posts, most of them in Dutch, and one seven-part English series, *Is AI the bicycle of the mind?*, which number 15 grows out of. The books: *Liever (g)een training* with Boudewijn Overduin, *Kostbaar misverstand* with Vivian Heijnen, which this series takes a good deal of its thinking from.

None of that is reprinted here. It has been reworked: brought up to date, tested against the research that has appeared since, and dropped where the evidence no longer carries it. Theory that has never met a client is a position; practice that cannot account for itself is a story. This series is what remains when each is held to the other. Performance Consultant AI carries the same discipline to the desk where the decision is actually made.

**About the cases.** Every worked case in this series comes from our own practice. Names are changed and details altered so that no organization is recognizable, and the arithmetic reconciles. That is said once, here, rather than under each case: client work is confidential, and a case that named anybody would be a breach rather than a proof.

THE SERIES

# Sixteen papers, *six parts*

Number 01 is the map and number 02 is the foundation. From number 05 on, every number is depth on exactly one step of the method. The filled block is the paper you are holding; the two unnumbered pieces close the set.

### I · WHAT HPT IS

*The discipline, where it came from, and why the reflex it replaces survives.*

- 01 The performance improvement process
- 02 Human performance technology
- **03 Environment before capability**
- 04 The system and the performer

### II · THE BUSINESS ANALYSIS

*Question 1.*

- 05 Business needs and performance needs

### III · THE PERFORMANCE ANALYSIS

*Questions 2 and 3.*

- 06 Gap analysis
- 07 The exemplary performer

### IV · THE CAUSE ANALYSIS

*Questions 4 and 5.*

- 08 Cause analysis: the three levels
- 09 Cause analysis: the six conditions

### V · THE DECISION

*From cause to solution, and what counts as proof.*

- 10 Intervention selection
- 11 Learning metrics and business metrics

### VI · THE FUNCTION

*The request, the money, the implementation, the assistant and the value.*

- 12 The solution request
- 13 The cost of non-performance
- 14 Implementation and maintenance
- 15 Performance augmentation
- 16 The Value Flip Calculator

**Arets L&D Business Models** · the destination piece, after 16
*No number: the four positions a function can hold, and the move the sixteen papers make simple.*

**Solutions that cannot work** · the field guide, at the end
*No number: it points at all sixteen. The same material, ordered by the sentence that lands on your desk.*

Business problem: a result is moving the wrong way · Training request: a solution has already been chosen · Both doors: after the decision

REFERENCES

# Where every claim *in this paper comes from*


Albarracín, D., Fayaz-Farkhad, B., & Granados Samayoa, J. A. (2024). Determinants of behaviour and their efficacy as targets of behavioural change interventions. *Nature Reviews Psychology, 3*(6), 377–392. https://doi.org/10.1038/s44159-024-00305-0

Garavan, T. N., McCarthy, A., Lai, Y., Clarke, N., Carbery, R., Gubbins, C., Sheehan, M., & Saunders, M. N. K. (2021). Putting the system back into training and firm performance research: A review and research agenda. *Human Resource Management Journal, 31*(4), 870–903. https://doi.org/10.1111/1748-8583.12337

Blume, B. D., Ford, J. K., Baldwin, T. T., & Huang, J. L. (2010). Transfer of training: A meta-analytic review. *Journal of Management, 36*(4), 1065–1105. https://doi.org/10.1177/0149206309352880

Gilbert, T. F. (1978). *Human competence: Engineering worthy performance.* McGraw-Hill. (The passage quoted in this paper is taken from the primary text as reproduced, with its page, in Winter, 2015, p. 29.)

Harless, J. H. (1985). Invited address: Performance technology and other popular myths. *Performance + Instruction, 24*(6), 4–6. https://doi.org/10.1002/pfi.4150240603

Hughes, A. M., Zajac, S., Woods, A. L., & Salas, E. (2020). The role of work environment in training sustainment: A meta-analysis. *Human Factors, 62*(1), 166–183. https://doi.org/10.1177/0018720819845988

Rummler, G. A., & Brache, A. P. (2013). *Improving performance: How to manage the white space on the organization chart* (3rd ed.). Jossey-Bass. (Original work published 1990)

Winter, B. (2015). *Agile performance improvement: The new synergy of Agile and human performance technology.* Apress.


*The three books named earlier in this paper are listed in full in number 02, so that they do not recur in every reference list.*