TULSER

No. 09 · © 2026 Tulser

WHITE PAPERS · PERFORMANCE CONSULTING

BUSINESS PROBLEM · QUESTION 5

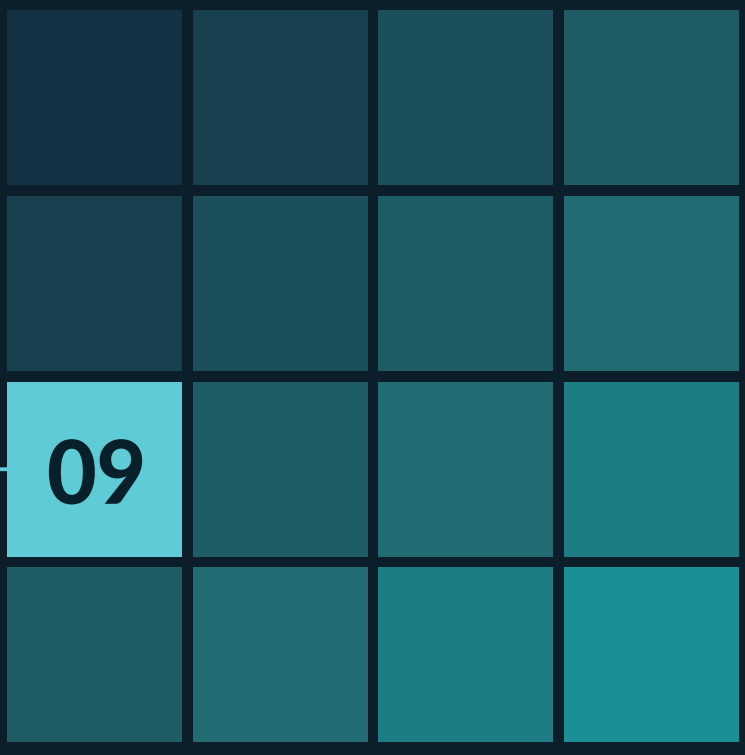
The level is established. Now: which of six conditions is actually binding?

Cause analysis: *the six conditions*

Six conditions, three of them arranged around the performer and three carried by them, how to establish which one is binding rather than which one is familiar, and what four decades of research since actually support for each.

JOS ARETS

with Vivian Heijnen



IN THIS PAPER

Only one of the six *lives inside the person.*

THE QUESTION — *The level is established and the gap is real. Which condition is actually preventing the work, and how do you find out, given that everybody in the room already has an opinion and only one of the six is visible from a desk?*

FOR — Anyone who owns a result or is asked to fix one: line managers and directors, operations, quality and continuous improvement, HSE, OD, HR business partners and L&D.

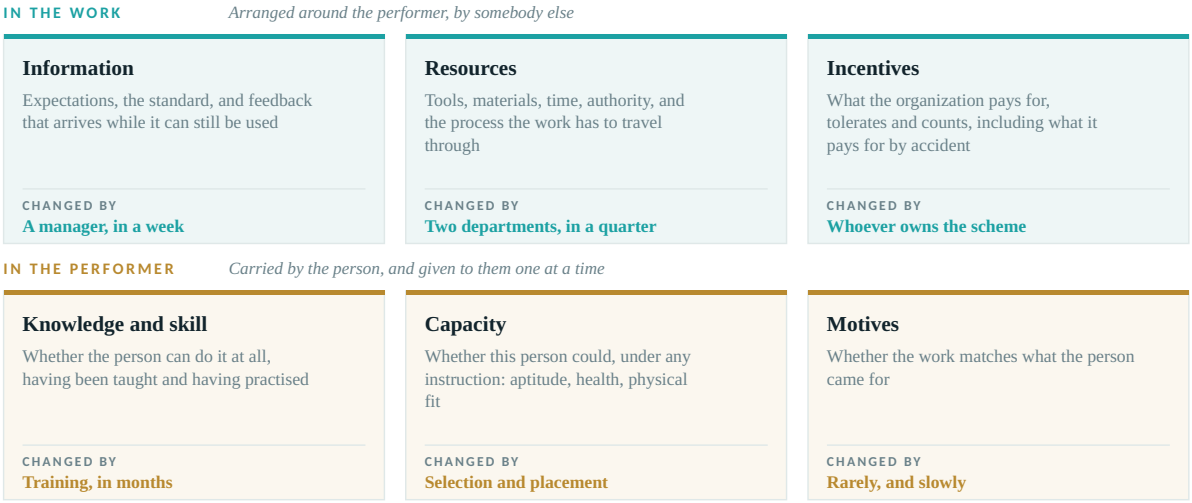
READ — 26 minutes · with a worked case and a one-page instrument.

01	Six conditions, and where each of them lives	Three are arranged around the performer. Three are carried by them.
02	What the model is, and what it is not	A checklist rather than a validated model, and why that is the right thing for it to be.
03	Condition by condition, what the research supports	Four decades of evidence, and the two conditions where it is genuinely strong.
04	Four questions that identify a binding condition	The order matters more than the questions do.
05	Rarely one, and rarely all six	Two failure modes: the single tidy cause, and the package that hides which part worked.
06	The dispatchers who all knew the rule	A worked case in which two conditions turned out to be binding, and skill was neither.
07	What this does not claim	Four limits, because the rest is not credible without them.
.	Run it yourself	Six conditions, on the gap in front of you.

Five of the six conditions are arranged by somebody else.
Only one of them is the person's to carry.

Six conditions,
and where each of them lives

Gilbert's contribution was not the list. It was the split. Three of the six conditions are arranged around the performer by somebody else, and three are carried by the performer, and the two halves are bought from different budgets, on different timescales, by different people.



Gilbert's six, in the arrangement that matters: what is arranged around the person, and what the person carries. The bottom row is where almost every proposed remedy lands.

What each one *looks like when it is missing*

A condition is not recognized by its name. It is recognized by what an organization looks like while it is absent, and those six pictures are worth more in a first conversation than the model is.

THE CONDITION	WHAT ITS ABSENCE LOOKS LIKE ON THE FLOOR	WHO CAN ARRANGE IT
Information: expectations and feedback	Capable people drifting, each with a private idea of what good is. Nobody finds out they were wrong until it is too late to act.	The manager who owns the result, usually within a week.
Resources: tools, time, materials	Workarounds. Second spreadsheets. Steps done twice. The data arrives after the decision it was meant to inform.	A system owner, a budget holder, and increasingly the team itself, for the part it controls.
Incentives	The wanted behavior is the punished one. What gets discussed at the meeting points away from the result being asked for.	Whoever sets what is reported and rewarded. One decision, and often free.
Knowledge and skill	Somebody genuinely does not know how, and would not be able to do it if their life depended on it.	A designer and a supplier. This is the one a course reaches.
Capacity	The right task in the wrong hands, or too few hands for the work as designed. Effort is not the problem and more of it will not help.	Recruitment, placement, or a redesign that removes the requirement.
Motives	“They just don’t seem to care”: said about a job with no autonomy, no visible result and no say in how the work is done. The complaint names the person; the design made the job.	Whoever designs the job. Work design rather than a pep talk, and section 03 says how thin the evidence under this row runs.

Six conditions, after Gilbert (1978) in the ordering proposed by Chevalier (2003). Read the middle column first: it is what a practitioner actually meets, and the name follows from it.

The asymmetry
the split reveals

Set the six out that way and something becomes hard to unsee. The top row is cheap, fast and owned by someone with a name. The bottom row is slow, expensive and owned by nobody in particular, and the bottom row is what the organization buys.

	THE THREE IN THE WORK	THE THREE IN THE PERFORMER
Who arranges it	A manager, a process owner, a system, a scheme. Someone whose name is on it.	Nobody, exactly. Recruitment did part of it, training did part, and the rest arrived with the person.
How long it takes	A week for information, a quarter for a process, one decision for an incentive.	Months for skill, a hiring cycle for capacity, and for motives, redesign the job rather than the person (section 03).
What it costs	Usually attention rather than budget. Often nothing but a decision.	A supplier, a room, working hours away from the job, and a per-head figure.
How visible it is	Invisible until somebody goes and looks at the work being done.	Entirely visible. It has a face, an appraisal and a training history.

Four rows, one pattern. The reflex described in number 04 is not irrational: it is the last column being the only one anybody can see.

> **WHY THIS QUESTION SITS LAST**
>
> Because it is the only one that cannot be answered from data. The four before it can be settled with records: a measure, two values, a distribution, an event log. This one requires somebody to watch the work and ask the performer questions in the right order. Asked first, it is a guess with a vocabulary. Asked fifth, the field has been narrowed to one level and one group of people, and the questions become answerable in an afternoon.

One thing the split does not claim. It does not say the environment is always the cause. It says the environment is the part you have not looked at, that it is the cheaper half to change, and that the sixth condition is the one everybody proposes before anybody has examined the other five. Number 03 gives the evidence for that ordering; this paper gives the method.

What the model is, *and what it is not*

Honesty first, because this is the point at which a skeptical reader is entitled to ask what the six conditions actually rest on. The answer is straightforward and worth stating plainly.

Gilbert's six are not a validated instrument and this series does not present them as one. There is no meta-analysis of the model, and there is unlikely ever to be one, for the reason number 02 sets out: a method whose purpose is to establish that *this* problem in *this* organization has *this* cause does not generate a population of comparable cases for anybody to pool.

What the six are is a diagnostic checklist, and a checklist earns its place differently from a theory. It earns it by being exhaustive enough that nothing important falls between its categories, and ordered well enough that the cheap questions come before the expensive ones. On both counts it holds up, and the individual categories it names have accumulated substantial synthesized evidence around them, built by researchers who had never heard of Gilbert. Section 03 sets that out condition by condition, including where it is thin.

> **WHICH IS THE HONEST BASIS FOR USING THEM**
>
> Not "the model has been proven", which would be false. Rather: the six categories carve the problem at joints that four separate research literatures also found, and using them costs an afternoon while getting the diagnosis wrong costs a program. That is weaker than a validation study and a great deal stronger than a preference.

One further consequence follows from the split, and it is the reason the checklist is worth its afternoon. Gilbert drew the line at three and three, but look at who can arrange each. Training can be bought, and capacity can be selected for or designed away, which leaves only motives truly carried in person: five of the six conditions can be arranged around the performer by somebody else. A practitioner who does not have a list will examine the sixth anyway, because the sixth is the one with a face, which is the whole subject of number 04.

We spend too much of our time "fixing" people who are not broken, and not enough time fixing organization systems that are broken.

GEARY A. RUMMLER & ALAN P. BRACHE

Improving Performance, 3rd ed., 2013, p. 11

Condition by condition, *what the research supports*

Four decades of work has accumulated around these six categories without ever addressing the model itself. Read across it and the strength of the evidence is markedly uneven, which is itself useful, because it tells you which findings will survive a skeptical room.

CONDITION	WHAT THE EVIDENCE ESTABLISHES	HOW STRONG
Information expectations and feedback	Feedback is not a generic lever. Across 435 studies and some 61,000 people, feedback helped on average, and about one time in six it made performance worse. What separates the two is how much the feedback actually tells you: feedback that says what the task requires, how the work is going and what to adjust does roughly four times as much as praise or reprimand (Wisniewski et al., 2020). The older workplace synthesis found the same, with a larger failure share (Kluger & DeNisi, 1996).	**Strong**, and inconvenient. The most-prescribed remedy in this row fails whenever it carries no information about the work.
Resources tools, time, authority, process	A systematic review of 55 work-redesign interventions found 39 reporting a positive effect on performance, two negative and 14 mixed, with implementation quality and alignment of surrounding systems as named boundary conditions (Knight & Parker, 2021).	**Good.** Note the boundary condition: a redesign that the surrounding systems contradict does not hold.
Incentives consequences, deliberate and accidental	Paying for performance does raise it, modestly, across 108 samples and some 71,000 people, but by two paths that pull against each other: it raises interest in the work and it raises pressure, and the pressure costs back part of what the interest gains (Chen et al., 2023). And across 183 studies and 212,000 people, incentives predicted *how much* got done while interest in the work predicted *how well*, and tying pay directly to output made the second matter less (Cerasoli et al., 2014).	**Strong, and two-sided.** The net effect is modest and its sign depends on the design.

The three conditions in the work. The middle column is what you can say in front of a finance director without being asked for a source you do not have.

And the same,
in the numbers themselves

The table says what the research supports. This says it in the form the research reports it, for the reader who wants to check rather than take it on trust.

THE FIGURES, FOR ANYONE WHO WANTS THEM

Feedback. Wisniewski and colleagues: 994 effect sizes from 435 studies, ~61,000 subjects; overall d = .48; about 17 percent of effects negative; high-information feedback d = .99 against d = .24 for reinforcement and punishment.

Work redesign. Knight and Parker: 55 interventions, 39 positive, two negative, 14 mixed, with implementation quality and alignment of surrounding systems as boundary conditions.

Incentives. Chen and colleagues: pay for performance relates to job performance at ρ = .23 across 108 samples and 71,438 people, through two opposing paths. Cerasoli and colleagues: 183 studies, 212,468 people; incentives predicted quantity, intrinsic motivation predicted quality.

And the three
the person carries

The bottom row is where the money goes, and it is where the evidence is thinnest, not because the conditions are unreal, but because they are far harder to change than the budget implies.

CONDITION	WHAT THE EVIDENCE ESTABLISHES	HOW STRONG
Knowledge and skill	Number 03 carries this evidence in full: across domains, interventions that add knowledge move behavior least, and interventions that change what somebody can reach move it most (Albarracín et al., 2024). Where the cause genuinely is a skill shortfall, well-designed instruction is exactly right and nothing substitutes for it.	**Strong, and widely misread.** It is a claim about how often, not about whether.
Capacity	The classic validity figures for selection were overstated: correcting a systematic overcorrection for range restriction lowers estimates by roughly .10 to .20, with structured interviews ranking highest (Sackett et al., 2022). And capacity is not one thing: narrow abilities least correlated with general ability add substantial incremental validity for task and training performance (Nye et al., 2022).	**Moderate.** Selecting for capacity works, considerably less well than the field long claimed.
Motives	Little that is usable at the level of an individual diagnosis. What synthesized evidence exists sits in the work-design literature, and it is about the job rather than the person, which relocates most of this condition into the top row.	**Weak, as a cause you can act on.** Treat a motives finding as a hypothesis, not a conclusion.

Three conditions, and the third one is where analyses go to die. “They are not motivated” is the most common finding in this row and the hardest to act on. *[practitioner observation, not a measured frequency]*

Four questions
that identify a binding condition

Mager and Pipe turned this into a decision sequence for practitioners, and the sequence is still the shortest statement of the method (Mager & Pipe, 1997). Four questions, and the order carries most of the value.

ASK	HOW TO ASK IT	WHAT A YES MEANS
1 · Does the performer know what is expected, and how they are doing?	Not "were they told": ask them to state the standard in a number, and ask when they last saw their own figure. Then check whether the figure they saw was accurate and recent enough to act on.	The first condition is binding, and the remedy is usually a report and a conversation rather than a program. This is the cheapest finding in the method and the most frequently missed.
2 · Could they do it today if everything depended on it?	The old question, and it still works. If the answer is yes, it is not knowledge and skill: whatever else it is.	A no points at the sixth condition or at capacity. A yes eliminates the most expensive remedy on the table in about four minutes.
3 · Does anything in the work prevent it?	Watch it being done. Look for the second spreadsheet, the wait for an authorization, the tool that fights the task, the information that arrives a day late.	The second condition. Almost always visible to the performer and almost never in any document, which is why this question has to be asked where the work happens.
4 · What happens to the person who does it right?	And what happens to the one who does not? Include what happens accidentally: more work as a reward for speed, a queue that punishes the careful, a target that pays for volume.	The third condition. The accidental consequence is the common finding here, and it is rarely anybody's intention.

Four questions, in this order, because each one removes remedies from the table before the next is asked.

Two rules about how they are asked. Ask them of the performer, not about the performer: the person doing the work knows which of the four is true and is rarely asked. And ask the exemplary performer the same four, because the difference between the two answers is the finding, which is the subtraction number 07 sets out.

Rarely one,
and rarely all six

Two failure modes sit on either side of this question, and an analysis can walk into either one on the same afternoon.

The first is naming a single cause where three are operating. It happens because the first condition that produces a clear answer is satisfying, and because a single cause makes a tidier report. The correction is arithmetic rather than judgement: for each condition you name, say what share of the gap it plausibly accounts for, and check that the shares are of the right order. Where they add to a fraction of the gap, something is missing and the analysis is not finished.

The second is the over-correction. Having found three conditions, propose something for each, and present the package as thoroughness. It reads well and it is expensive, and it destroys the one thing the analysis was for: after a package, nobody can say which part worked. That is not merely an evaluation problem for later: it is what makes the next analysis harder, because the organization now has a success story with no mechanism in it.

> **THE WORKING RULE**
>
> Name every condition the evidence supports, with what supports it, and rank them by the share of the gap each accounts for. Then act on the binding one first and re-measure. Number 10 takes the choice of remedy: one per established condition, and why that is not the same as buying several remedies at once; number 11 takes what may be claimed afterwards.

One practical note about ranking. The share a condition accounts for and the cost of reaching it are different axes, and the binding condition is not automatically the one to act on first. A condition that explains a fifth of the gap and can be fixed by changing a report next week will usually be done before one that explains half and needs two directors to agree. Say both numbers and let the person who owns the result decide, which is the limit number 02 sets on the whole discipline.

Everybody knew the rule.
Nobody was ever told they had broken it.

The dispatchers
who all knew the rule

A national maintenance contractor, 62 dispatchers across three centers. The measure was the share of jobs dispatched with a complete parts list. Desired 90 percent; running at 61 over four quarters. The proposal on the table was a refresher on the dispatch procedure.

Question four had established the level: uniform across all three centers and all 62 dispatchers, with nothing concentrated anywhere, and the ceiling test showed that even at the exemplary performer's rate the unit would reach 90. So the cause was reachable by people, but which condition?

QUESTION	WHAT THE ANSWER WAS	WHAT IT ELIMINATED
1 · Do they know what is expected, and how they are doing?	Every dispatcher named the 90 percent target. Not one of them had ever seen their own figure. The measure existed in a monthly report that went to the three center managers and stopped there.	Nothing yet, but this is the first candidate, and it took eleven minutes to establish.
2 · Could they do it today if everything depended on it?	Yes. Asked to complete a parts list for a job in front of them, all six dispatchers in the sample did it correctly and unhurriedly.	Knowledge and skill. The refresher course, which was the entire proposal, was off the table before lunch.
3 · Does anything in the work prevent it?	Partly. The parts catalogue took four clicks and about ninety seconds per line to search. Irritating, and not prohibitive.	Nothing, but it produced a second finding worth recording.
4 · What happens to the person who does it right?	They take longer on the call. Call handling time was reported weekly, per dispatcher, and discussed in every team meeting. Parts-list completeness was reported to nobody.	This is the finding. The organization measured and discussed one thing and wanted another.

Four questions, half a day. Two conditions were binding and the second one was paying for the behavior it did not want.

What was changed, *and what it cost*

Two remedies, both in the top row of the six, and neither of them a program.

The first: the completeness figure was added to the weekly report that already existed, per dispatcher, next to call handling time. That is a report change, and it took an analyst two days. The second: the team meeting agenda was changed so that the two figures were discussed together rather than one of them alone, which is a five-minute decision by a center manager, and it was by far the harder of the two to get.

MEASURE	BEFORE	AFTER TWO QUARTERS
Jobs dispatched with a complete parts list	61%	84%
Against the desired level of 90%	29 points short	6 points short
Average call handling time	4.9 minutes	5.4 minutes

The third row is the honest one. The remedy cost half a minute per call, which the organization had been unknowingly buying performance with all along.

Two notes about the ending. The remaining six points were the catalogue: the third question's finding, quantified afterwards at roughly the right size, and fixed the following year when the catalogue was replaced for unrelated reasons. And the analysis had to say out loud that handling time would rise, before the change, because a remedy that quietly moves a second number and does not warn anybody is how a good finding loses its credibility in the following quarter.

What this *does not claim*

This is the paper in the series with the weakest formal foundation, and saying so is the condition for the rest being usable.

- **The six conditions are a checklist, not a validated model.** There is no synthesized evidence for the model itself and there is unlikely to be. What the six have is exhaustiveness, a useful order, and convergent support from literatures built independently of them.
- **The supporting evidence is about categories, not about the model.** Wisniewski and colleagues studied feedback in education, Kluger and DeNisi studied it at work and in the laboratory, and Knight and Parker studied work redesign. None of them studied Gilbert's conditions. Reading those syntheses as validation of the model is the error this paper is trying not to make; what they establish is that the category is real and that its effect is conditional.
- **Nothing here establishes how often each condition binds.** No study reports that distribution, for the same reason set out in number 08. The claim is about the order of examination, which holds whatever the distribution is.
- **The sixth condition is real.** Sometimes people genuinely cannot do the work and instruction is exactly the right answer. This paper argues that it is the last of six to examine and the first of six to be proposed, which is a claim about sequence and not about worth.

SUMMARY

Six conditions, five of them arranged around the performer by somebody else. The evidence for the individual categories is real and uneven, and strongest exactly where organizations look least: at feedback that never arrives and at consequences nobody intended. The method is four questions in a fixed order, asked of the person doing the work, and its whole value is that each one removes a remedy from the table before the next is asked.

RUN IT YOURSELF

Six conditions, *on the gap in front of you*

Half a day, with the people doing the work rather than with the people describing it. Line one is answered in the room; the rest are answered where the work happens.

1 · The standard, in a number

Ask three performers to state it. Do they agree with each other, and with you? Then: when did each of them last see their own figure?

2 · Could they do it today

If everything depended on it. Ask them to do it in front of you. A yes takes training off the table and takes four minutes.

3 · What is in the way

Watch one person do the work for an hour. Write down every workaround. Then ask them what they would change if it were theirs to change.

4 · What happens to the person who does it right

And to the one who does not. Include the accidental consequences: more work, a longer call, a queue that punishes care.

5 · The exemplary performer, same four questions

Where do the two sets of answers differ? That difference is your finding, not your impression of the person.

6 · Rank what you found

Which condition accounts for the largest share of the gap? Act on that one and re-measure, not on all of them at once.

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

Cerasoli, C. P., Nicklin, J. M., & Ford, M. T. (2014). Intrinsic motivation and extrinsic incentives jointly predict performance: A 40-year meta-analysis. *Psychological Bulletin, 140*(4), 980–1008. https://doi.org/10.1037/a0035661

Chen, Y., Zhang, Z., Zhou, J., Liu, C., Zhang, X., & Yu, T. (2023). A cognitive evaluation and equity-based perspective of pay for performance on job performance: A meta-analysis and path model. *Frontiers in Psychology, 13*, 1039375. https://doi.org/10.3389/fpsyg.2022.1039375

Chevalier, R. (2003). Updating the behavior engineering model. *Performance Improvement, 42*(5), 8–14. https://doi.org/10.1002/pfi.4930420504

Gilbert, T. F. (1978). *Human competence: Engineering worthy performance.* McGraw-Hill.

Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin, 119*(2), 254–284. https://doi.org/10.1037/0033-2909.119.2.254

Knight, C., & Parker, S. K. (2021). How work redesign interventions affect performance: An evidence-based model from a systematic review. *Human Relations, 74*(1), 69–104. https://doi.org/10.1177/0018726719865604

Mager, R. F., & Pipe, P. (1997). *Analyzing performance problems, or you really oughta wanna* (3rd ed.). Center for Effective Performance.

Nye, C. D., Ma, J., & Wee, S. (2022). Cognitive ability and job performance: Meta-analytic evidence for the validity of narrow cognitive abilities. *Journal of Business and Psychology, 37*(6), 1119–1139. https://doi.org/10.1007/s10869-022-09796-1

Sackett, P. R., Zhang, C., Berry, C. M., & Lievens, F. (2022). Revisiting meta-analytic estimates of validity in personnel selection: Addressing systematic overcorrection for restriction of range. *Journal of Applied Psychology, 107*(11), 2040–2068. https://doi.org/10.1037/apl0000994

Wisniewski, B., Zierer, K., & Hattie, J. (2020). The power of feedback revisited: A meta-analysis of educational feedback research. *Frontiers in Psychology, 10*, 3087. https://doi.org/10.3389/fpsyg.2019.03087

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.