TULSER

No. 04 · © 2026 Tulser

WHITE PAPERS · PERFORMANCE CONSULTING

BUSINESS PROBLEM

Everyone can see the people. Almost nobody can see the work.

The system and *the performer*

How a system stays invisible while the people inside it stay in plain view, what that costs, and the five moves that make the work show up again.

04

JOS ARETS

with Vivian Heijnen

IN THIS PAPER

The work
has no face.

THE QUESTION — *A result is disappointing and every explanation on the table is about people: ownership, attitude, capability. Why does the conversation keep landing there, and what would it take to see the work instead?*

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

01	The asymmetry	A process has no face. The people in it do.
02	Four reasons it survives	Why the reflex outlives every attempt to argue it away.
03	What the research calls it	Fifty years of attribution work, and what it does and does not license.
04	The space between the boxes	Three levels of performance, and why the middle one is invisible.
05	The problem with a person's name	A worked case in which the supervisor was the only visible part of the system.
06	Five moves that make the work visible	None of them takes longer than a week.
07	A case in which we got it wrong	Our own analysis: correct, incomplete, and wrong for a reason this paper caused.
08	What this does not claim	Three limits, because the rest is not credible without them.
·	Run it yourself	What can you actually see, on the problem in front of you?

The work has no face.
The people in it do.

A process *has no face*

If the evidence has been clear for half a century, the obvious question is why every organization still reaches for the person first. The answer is not that managers are careless. It is that one of the two candidate causes in front of them is visible and the other is not.

Walk a floor and you can see people. You can see who arrives early, who looks harried, who asks good questions and who avoids the difficult customer. You know their names, you have appraised them, you have an impression of each of them that took years to form and takes seconds to recall.

Now try to see the process. It has no location. It crosses four departments, three systems and two shifts. Nobody is accountable for the whole of it, so nobody reports on the whole of it. Its worst moments happen in the gaps between the people you can see: in the two days an order sits in a queue that belongs to no one. There is no meeting at which it is present.

This asymmetry decides where the conversation goes. When a result disappoints, the explanations that come to mind are made of the material that was available to the mind, and what was available was people. Nobody chose to look in the wrong place. They looked in the only place that could be looked at.

> **WHY THIS MATTERS MORE THAN IT SOUNDS**
>
> Every paper after this one describes a method for examining the system. None of it will be used by someone who is confident the problem is a person, because there is nothing to examine. The habit of looking at the work is not a technique. It is the precondition for the technique.

Four reasons
the reflex outlives the argument

Being told that the environment holds the larger gain does not change what a manager does on Monday. Four forces hold the reflex in place, and each of them is rational on its own terms.

WHAT HOLDS IT	WHY IT IS REASONABLE	WHAT IT PRODUCES
The person is observed; the system is inferred	One is evidence of the senses, the other is a conclusion that has to be constructed from records.	The explanation that arrives first is about somebody.
The person can be addressed today	A conversation takes twenty minutes. Changing a handoff takes three departments and a quarter.	The available action displaces the effective one.
The language is built for it	Appraisals, competency frameworks, talent reviews and training budgets all take the individual as the unit.	There is no form on which to record “the process did this”.
The confirming case	Sometimes it genuinely is the person, and that instance is remembered vividly and cited for years.	A rule of thumb kept alive by its rarest success.

Four reasons, none of them foolish. Which is why the reflex cannot be argued away and has to be worked around instead.

The fourth is the most durable, because it is true. People do sometimes cause problems, and when they do, addressing them works quickly and unmistakably. The trouble is that a rule which is right one time in six will feel right every time it is right, and feel like bad luck the other five. Nothing in the ordinary experience of managing corrects it.

There is a fifth reason that belongs to the organization rather than the individual: the people closest to the problem often cannot say what they work around. Reporting that a system forces a daily improvisation is, in many organizations, indistinguishable from admitting you have been doing it wrong. Where speaking up is expensive, the system stays invisible by consent. Pooling 136 independent samples, psychological safety relates to information sharing and to performance across sectors and countries (Frazier et al., 2017); the review of where that literature now stands is Edmondson and Bransby (2023).

> *An organization behaves as a system, regardless of whether it is being managed as a system. If an organization is not being managed as a system, it is not being effectively managed.*

GEARY A. RUMMLER & ALAN P. BRACHE

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

What the research *calls it*

This is one of the most thoroughly studied findings in social psychology, and it has a name. Knowing the name is useful mainly because it stops the conversation being about anyone's competence.

When people explain someone else's behavior, they systematically overweight what the person is like and underweight the situation the person was in. The effect was named half a century ago and has been examined ever since. Reviewing what had accumulated, Daniel Gilbert and Patrick Malone gathered it under the more careful label *correspondence bias* and established something more useful than the effect itself: it is strongest precisely when the situation is hard to observe (Gilbert & Malone, 1995). The name is a coincidence the field has learned to live with; this is not Thomas Gilbert, whose six conditions run through the rest of this series.

That last finding is the one that matters here. The bias is not a general failure to think about context. It is a failure that grows as the context becomes harder to see, which is an exact description of a cross-functional process.

> **WHAT THIS DOES AND DOES NOT LICENSE**
>
> It licenses one conclusion: an explanation that arrives quickly and is about a person deserves a second look, because that is the explanation the mind produces by default. It does not license the opposite reflex. "It is always the system" is the same error wearing different clothes, and it is just as unfalsifiable. The point of the analysis is to establish which, on evidence, in this case.

How far this evidence *actually reaches*

It is worth being exact about how far this evidence reaches, because it does not reach all the way by itself. The attribution studies asked people to explain the behavior of strangers: a student who had read out an essay, a driver who cut somebody up, a candidate in an interview. Nobody has ever sat a management team in front of a delivery report and measured how they explain it.

So the step from that research to your Tuesday meeting is an inference rather than a finding. What carries it is the mechanism, not the setting. In the laboratory and in the meeting the same asymmetry holds: one thing is there to be seen and the other has to be assembled out of records, and the bias grows exactly as the second becomes harder to assemble (Gilbert & Malone, 1995). By that measure a management meeting is close to the worst case there is. The people are in the room, with names and reputations. The process is in four systems and nobody has printed it.

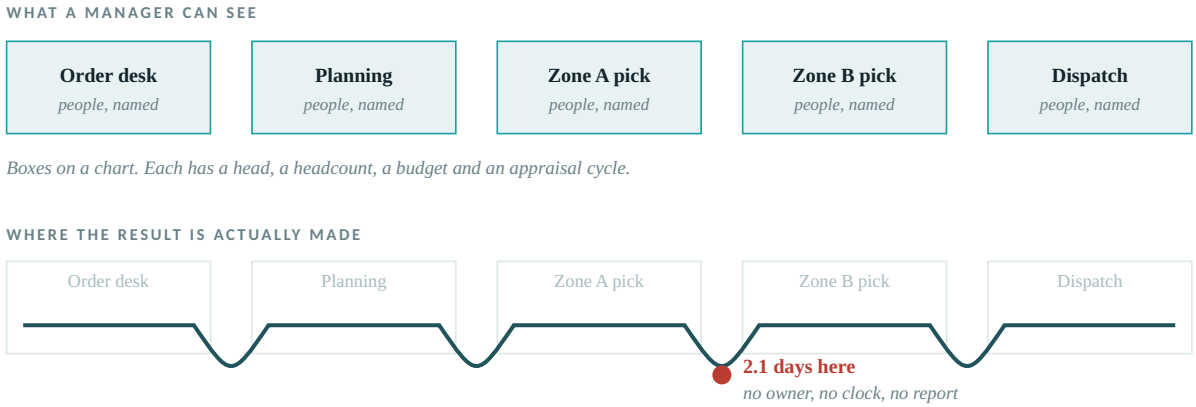
It is a well-supported inference: Rummler and Deming arrived at the same conclusion from inside organizations, by different routes and without the psychology. But it is an inference, and this paper says so rather than dressing it as a result.

> **WHAT WOULD SETTLE IT**
>
> A study that recorded how management teams explain a disappointing result, and then established independently where the cause actually sat, would turn the inference into a finding. As far as we can establish, nobody has run it. Until somebody does, the honest position is the one taken here: the mechanism is well evidenced, the setting is not, and the practical advice, look at the work before you look at the people, does not depend on closing that gap, because it is cheap to follow and easy to reverse if the work turns out to be fine.

The space
between the boxes

Rummler put the invisibility in one image that has outlasted every model built on it: the organization chart shows the boxes, and the work happens in the space between them.

Performance can be examined at three levels: the organization, the process that crosses it, and the person doing the work. The first is on every strategy page. The third is on every appraisal. The second, which is where most results are actually made or lost, is on nobody's report (Rummler & Brache, 2013).



The same five boxes, twice. The upper panel is what gets staffed, budgeted and appraised; the lower one is where the result is made.

Why the middle level *has no owner*

The gaps are invisible for a structural reason, not a cultural one. Every function measures its own throughput, and no function measures the wait between two functions.

Zone A reports what it picked. Zone B reports what it picked. Neither reports the two days an order spent between them, because that time belongs to neither and appears in neither system as work. It is visible only to the person waiting for the order and to the customer at the end of the line, and neither of them attends the operations meeting.

An example makes it concrete. Four departments each hold a one-day target for their own step, and each one meets it, week after week. That is four days of work. The promise to the customer is five days. The customer waits nine.

Nobody has done anything wrong and no report shows a problem. The five missing days were spent in the three spaces between the four departments, and a department does not report on a space it does not own. Ask each of the four heads whether their step is on target and every answer is yes, truthfully. Ask where the other five days went and there is nobody to ask, because the question belongs to no job.

That is the pattern, and it is recognizable once you have seen it: every department's numbers look acceptable while the end-to-end number is unacceptable, and the discussion turns to whether people are pulling together. Deming had the arithmetic for it long before anyone had the diagram: most of what goes wrong belongs to the system, and the system belongs to management (Deming, 2000, p. 270).

> **THE TELL**
>
> When each function is meeting its targets and the customer is not being served, you are looking at a white-space problem, and no amount of examining the functions will find it. The question that opens it is not "who is underperforming?" but "where does the work stop?".

The system is not hidden.
It is simply not on anyone's report.

The problem
with a person's name

A building-products distributor, 68,000 orders a year, five branches. The figures are illustrative and reconcile from the rates stated with them.

The promise to customers was delivery within five working days. Roughly three orders in ten were arriving late, and the explanation everyone offered was the same: the warehouse had an ownership problem, and specifically it had a supervisor with a reputation. Two branch managers had said so independently. A performance conversation was scheduled.

Nobody had split the number. When somebody did, orders drawn from a single picking zone were running at 4.7 days and orders that needed both zones at 11.2. The split ran straight through the supervisor everyone had named: he ran the multi-zone queue, and every order that arrived late arrived late through him.

WHAT THE SPLIT SHOWED	
Single-zone orders	4.7 days, comfortably inside the five-day promise. Seven in ten of the volume.
Two-zone orders	11.2 days, all of them late. Three in ten of the volume: 20,400 orders a year.
Where the time went	2.1 days waiting for the second zone's picks to be walked across. The transfer had no owner, no schedule and no timestamp, so it appeared in no system.
What it cost	At $84 (€74) an order in expedited freight and service credits, **$1.71 million (€1.51M) a year**.
Who already achieved it	One branch ran two-zone orders at 5.2 days. A shift lead there walked the second zone's picks over herself at eleven and at three, every day, and had done for a year without mentioning it.

The organization had a name for the problem, and it was a person's name.

What was visible, *and what was not*

Nothing about this case is subtle once it is written down. That is the point worth sitting with: it was not subtle before either, and it had run for a year.

What was visible: a supervisor, a queue that was always behind, and two branch managers with a consistent impression. What was not visible: a transfer between two zones that nobody had scheduled, nobody timed and nobody owned, and a shift lead in another branch who had quietly solved it and had no reason to think her walk to the far end of the warehouse was worth reporting.

The fix was a scheduled zone-to-zone transfer twice a shift and a timestamp on the handoff, so that the wait would appear in a system for the first time. The timestamp was a configuration change. The performance conversation did not take place.

> THE TWO HABITS THIS CASE REWARDS
>
> **Split the number by something nobody splits it by.** One cut, single-zone against two-zone, turned a personality into a process, and it took an afternoon. **Ask who already achieves it.** The answer was standing in another branch, and what she had was not a better attitude. It was a walk that somebody should have put on a schedule.

Both moves belong to the analysis proper: the split is question two and the exemplary performer is question three, taken up in numbers 06 and 07. They appear here only to make the point that the invisible becomes visible faster than anyone expects, once somebody looks.

Five moves that *make the work visible*

None of these is an analysis. They are what you do in the week before one, and any of them can be done by the person who owns the number, without a budget.

MOVE	WHAT IT SURFACES
1 · Follow one unit of work end to end	One order, one claim, one application: from the moment it arrives to the moment it is invoiced. Not a workshop about the process; the actual route of an actual item.
2 · Count the handoffs	Every point where the work changes hands, changes system, or waits for a decision. Each one is a candidate. Most organizations are surprised by the number.
3 · Split the number by something nobody splits it by	Single-line against multi-line, new against repeat, day shift against night, one region against the rest. A cut that has never been made is where the problem is hiding.
4 · Find who already achieves it	In our practice somebody almost always does, under the same conditions and with the same tools. What they have that others lack is evidence the organization already owns.
5 · Read the workarounds	A private spreadsheet, a laminated card, a WhatsApp group, a walk to the far end of the warehouse. Every one of them is a map of a system that does not work.

Five moves, a week at most. Together they turn the middle level of performance from an abstraction into something with a location and a duration.

The fifth deserves emphasis because it is the cheapest and the most consistently ignored. Nobody builds a private workaround for fun. Each one exists because the official system fails at a specific point, and the person who built it knows exactly where that point is. All that is required is to ask without making it an accusation, which brings the argument back to why the fifth reason in section 02 matters so much.

A case
in which we were the problem

Every case in this series so far ends well, which ought to make a reader suspicious. This one is ours and it went wrong: in the direction this paper is least comfortable with.

A field service operation, fourteen engineers, first-time fix rate 68 percent against a required 82. We ran the analysis exactly as this paper recommends. Nobody was blamed. We followed one job end to end and found a plausible systemic cause: the van stock list had not been revised after a product change, so engineers were arriving without the part in roughly three visits out of ten. The list was rebuilt. It was cheap, it was real, and we wrote it up as the finding.

Six weeks later first-time fix stood at 71 percent. Moved, and nowhere near 82. The service manager rang to ask what we had missed. We did not know, which is the answer no consultant wants to give, so we went back and looked again.

WHAT THE SECOND LOOK FOUND

What the van stock actually explained	Three points of the fourteen. Our finding was correct and it was small.
Where the remaining eleven sat	In two engineers of the fourteen. Twelve were running at 79 percent; two were at 23.
Why	Both had joined after the last training round on the new controller and had been picking it up from colleagues. Neither had ever been certified on it.
What that is	The sixth condition. Knowledge and skill, genuinely absent, in two named people.

$0.86 \times 79 + 0.14 \times 23 = 71$. The company-wide figure was an average of two situations that had nothing to do with each other.

The two were certified. First-time fix reached 78.3 percent and stopped there, still short of the 82, and the last four points we never accounted for. That is where the case ends, because that is where it actually ended.

Why we *missed it*

The interesting part is not that training turned out to be part of the answer. It is why we, who knew this method and had written about it, did not see two engineers at twenty-three per cent.

We had cut the gap by region and by product type. Both cuts came back uniform (every region a little short, every product type a little short) which is the pattern that says *look at what they all share*. So we looked at what everybody shared, found something real, and stopped. We never made the cut by engineer, and at that cut the same data was violently concentrated. The records to make it had been in the first extract all along.

TWO LESSONS, AND THE SECOND IS ABOUT US

The first is number 06's rule, learned the expensive way: cut the gap at least three ways before concluding anything from its pattern. Two cuts agreeing is not confirmation; it can simply mean the same cut has been made twice.

The second is uncomfortable, because it is about us. We had accepted the argument in this paper and turned it into a habit. The person is the last place to look, so we looked at the person last, found something systemic on the way, and never got there. The reflex this paper attacks had been replaced by its mirror image, and we walked into it holding our own instrument.

The claim, *in its exact form*

So the claim in this paper needs stating one more time in its exact form. It is not that causes are environmental. It is that the environment is *harder to see*, and therefore has to be examined first, because the visible explanation will still be there afterwards. Examining it first is not the same as accepting it. An analysis that finds a systemic cause and stops has done half the job, and the half it skipped is the half everybody else does.

AND THE AUTHORITIES THAT DO NOT SUPPORT THIS SENTENCE

Three of them are quoted at this discipline constantly. Deming's ninety-four percent. Rummler and Brache on the good performer who loses to a bad system. Gilbert on where the improvement potential sits. All three say something about *where causes tend to be*, and the sentence above deliberately refuses to say that. Borrowing their authority for it would be the exact over-reading this series spends sixteen numbers warning against.

The warrant is already in this paper, in section 03, and it is a finding about seeing rather than about causes: the pull towards explaining by the person grows precisely as the situation becomes harder to assemble out of records (Gilbert & Malone, 1995). That supports a claim about order, which is the only claim being made. An explanation that is easy to reach is not thereby more likely to be right. It is merely earlier, and unlike the other one, it will still be sitting there in a week.

What this *does not claim*

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

- **People do cause problems.** Sometimes the cause really is capacity, or knowledge, or a person in the wrong role. The claim here is about the order of examination, not about the outcome of it. Look at the work first because it is harder to see, not because the answer is known in advance.
- **The attribution evidence is imported.** That literature studied how people explain the behavior of others, largely outside organizations. Applying it to how a management team explains a delivery number is an inference, and it is stated as one throughout.
- **Making the work visible is not the same as fixing it.** A timestamp on a handoff reveals a wait. Removing the wait requires two departments to change something, and that is a different kind of work with a different failure rate.

SUMMARY

Nothing in this paper is anyone's fault, and that is exactly what makes it durable. The work is not concealed; it is simply not reported, and what is not reported cannot be discussed. Five ordinary moves, none of them taking more than a week, put it back in the room.

RUN IT YOURSELF

What can you *actually see?*

Ten minutes, on the problem in front of you now. Answer these five and you will know whether you are looking at a person or at a system. Leave one blank and you have found the first thing to establish.

1 · The explanation on the table, and whose name is in it

Write it exactly as it was said to you. If it contains a person, underline them.

2 · One unit of work, followed end to end

Which order, claim or application, arriving when, invoiced when? Trace it, do not describe it.

3 · The handoffs you counted

Every change of hands, system or decision-maker. How many, and how long does the work wait at each?

4 · The split nobody has made

Cut the number one way it has never been cut. Which cut, and what did it show?

5 · Who already achieves it, and what they do differently

Name a person, shift, site or unit, and the workaround they have never mentioned.

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*

Deming, W. E. (2000). *Out of the crisis* (2nd ed.). The MIT Press. (Original work published 1986)

Frazier, M. L., Fainshmidt, S., Klinger, R. L., Pezeshkan, A., & Vracheva, V. (2017). Psychological safety: A meta-analytic review and extension. *Personnel Psychology, 70*(1), 113–165. https://doi.org/10.1111/peps.12183

Edmondson, A. C., & Bransby, D. P. (2023). Psychological safety comes of age: Observed themes in an established literature. *Annual Review of Organizational Psychology and Organizational Behavior, 10*, 55–78. https://doi.org/10.1146/annurev-orgpsych-120920-055217

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117*(1), 21–38. https://doi.org/10.1037/0033-2909.117.1.21 (Daniel T. Gilbert, no relation to Thomas F. Gilbert.)

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

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)

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.