DR. IMME BAUMÜLLER DIGITAL LEADERSHIP ADVISORY

Embed

If Going Back Is Easy, People Go Back

The human side of AI transformation

PART 4 OF 4

Embed: If Going Back Is Easy, People Go Back

You can Engage and Enable a team to perfection and still watch them revert the moment attention moves on. Lasting behaviour change is not a training outcome — it is a design outcome. This essay is about the two levers that decide whether the new way sticks: the structures around people, and the routines inside them.

PART 4 OF 4 · A SHORT SERIES ON THE HUMAN SIDE OF AI TRANSFORMATION

> A quick orientation: Part 1, *The AI Paradox*, asked why AI investments stall — most companies end up pilot-rich but transformation-poor. Part 2, *Engage*, was about turning a vision into a direction people can act on. Part 3, *Enable*, was about giving them the ability to act. This part, *Embed*, is about making the new behaviour the default — so it survives the week the pilot ends.

THE KEY TAKEAWAYS, UP FRONT

- One useful way to look at Embed is on two levels. **Macro-Embed** is structural — processes, org design, governance, rewritten role profiles, old steps eliminated not just automated. **Micro-Embed** is human — the individual routines and the reward system around them. This essay focuses on the human half.
- Habits take far longer to form than people expect. The landmark study puts the median at about **66 days** to automaticity, with a range of 18 to 254.[1] Two months of effort is normal — not failure.
- One uniform fix rarely works. Katy Milkman identifies **seven distinct barriers** to change, and different people may be stuck at different ones[2] — so you shape the environment, not just the message.
- The reward system can quietly beat the training. **57% of employees** say they hide their AI use and present AI work as their own.[3] If saved time is quietly refilled with more work, people have a rational reason to conceal the gain.
- The numbers that decide Embed live in the performance sheet and the budget. The evidence says invest at least **one euro in people for every euro in technology** — 1:1 as a realistic floor, with some guidance running higher.[4]

EMBED · PART 4 OF 4

Two Levels of Embed

One useful way to approach the problem — the structures around people, and the routines inside them

STRUCTURAL Macro-Embed

- End-to-end process redesign
- Changes to organisational structure
- Governance rebuilt for hybrid human–agent work
- Role profiles rewritten around AI-augmented work
- Old process steps *eliminated*, not just automated

HUMAN Micro-Embed

- Individual habit formation — the new tool as default
- Incentives & performance management, redesigned
- Psychological safety for experimentation
- Knowledge made visible — gains turned collective

This essay goes deep on the human half.

A new behaviour laid on top of an unchanged structure rarely survives.

DIGITAL LEADERSHIP ADVISORY
The Human Side of AI Transformation

Two Levels of Embed — the structures around people, and the routines inside them.

Everyone reverts.

There is a moment every transformation lead knows and dreads. The pilot went well. People were trained, the tool worked, the early numbers looked good. Then attention moved on to the next initiative — and three months later the team is quietly back to the old routine. The templates they always used are still the default. The report still gets built the way it was always built. Nothing dramatic happened. People simply drifted back, because the old way was still the easy way.

This is Embed, the third E of the Transformation Triad, and it is where — in my experience — a great many transformations actually die. Not at Engage, not at Enable, but here, in the gap between *people can work the new way* and *people do work the new way, by default, when no one is watching*. The uncomfortable truth underneath it is simple: if going back is easy, people go back. Training alone rarely gets you out of that. You have to design for it.

What the research shows

It helps to look at Embed on two levels — not as a textbook definition, but as a way to approach the problem from two angles at once.

Macro-Embed is structural. It is end-to-end process redesign, changes to org structure, a governance model rebuilt for hybrid human-agent work, and — the step most often skipped — role profiles rewritten around AI-augmented work, with old process steps eliminated rather than politely automated. That work overlaps heavily with process and operating-model design, and it deserves more room than one essay can give it, so I'll return to Macro-Embed in its own right in a future piece. The point to hold onto here: a new behaviour laid on top of an unchanged structure rarely survives.

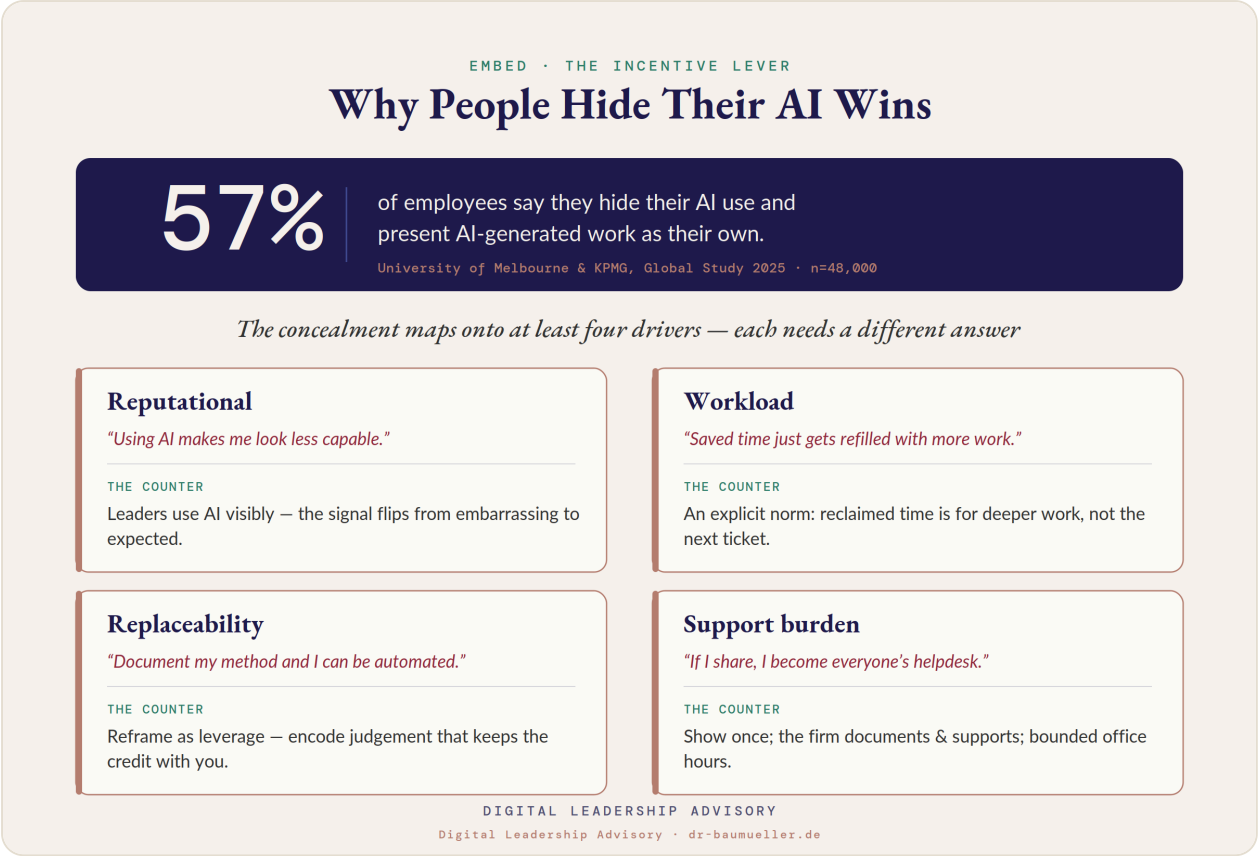
Micro-Embed is human — the individual routines, and the incentives that surround them. This is the half this essay goes deep on, through its two decisive levers: how habits actually form, and how the reward system either protects or quietly kills the new behaviour.

Start with habits, because the folklore here is wrong. The popular "21 days to form a habit" rule has no research behind it — it traces to a 1960s cosmetic-surgery book, an observation about patients adjusting to their appearance, never an experiment. The real figure comes from Phillippa Lally's landmark study: a median of about 66 days to reach automaticity, with a range from 18 to 254 days depending on the person and the behaviour.[1] A 2024 meta-analysis of twenty studies and roughly 2,600 people confirmed the shape of it — medians in the range of 59 to 66 days, means well over a hundred.[5] The lesson is not the number; it is the variation. If a new AI tool still feels effortful after three weeks, your people are not failing. They are on schedule.

What actually turns a step into a routine is well established: consistent repetition in a stable context, with the early repetitions mattering most; anchoring the new behaviour to a trigger that already exists ("after I open my inbox, I draft with AI first"); rewarding it immediately rather than at the end of the week; and — the finding that takes the pressure off — accepting that one missed day does not reset the progress.

And here is why a single roll-out rarely lands: Katy Milkman, a professor at the Wharton School, identifies seven distinct barriers to change, and different people may be stuck at different ones.[2] Somebody who struggles to get started might find a fresh-start moment helpful — a Monday, the first of the month, the first of May. Somebody who forgets is helped by a cue-based plan ("when X, then Y"). Somebody held back by low confidence often improves — counter-intuitively — when you let them give advice to others. Somebody who takes their cue from the room moves when they see visible peers doing it. Push one uniform programme at all seven and it may land for only a fraction of them. For you as a manager, the takeaway is this: you also need to consider the environment and shape it, not just sharpen the message.

Now the lever that, in my experience, carries a great deal of the weight: incentives. Often organisations invest in training and new workflows — which is exactly right — but some leave the performance-management system untouched, and that system can quietly win. Evidence for how invisible this might be comes from the University of Melbourne and KPMG's 2025 global study of 48,000 people across 47 countries: 57% of employees admit to hiding their AI use and presenting AI-generated work as their own.[3] Read that against the grain — people are concealing not their failures but their successes, the workflows that turned a three-hour task into twenty minutes.



Why people hide their AI wins — four drivers, each with its own counter.

Why hide a win? In my experience the concealment maps onto at least four drivers, and each needs a different answer:

- **Reputational** — "using AI makes me look less capable." The counter is visible use from the top: when leaders show their own AI use, the signal flips from embarrassing to expected.
- **Workload** — "if I finish in twenty minutes, I'll just be handed three more tasks; the saved time is taxed, not banked." The counter is an explicit norm that reclaimed time is for deeper work, learning or recovery — not automatically refilled.
- **Replaceability** — "once my method is documented in the system, it can be handed to someone cheaper, or automated entirely." The counter is to reframe the act as leverage rather than self-liquidation: you are encoding judgement that keeps the credit — and the role — with you.
- **Support burden** — "if I share, I become everyone's unpaid helpdesk." This one is easy to miss, and it stops many of your best people from sharing at all. The counter is to make

sharing cheap: the discoverer demonstrates once, the organisation does the documenting and the ongoing support, and bounded office hours replace the endless tap on the shoulder.

> A QUICK ASIDE — AND AN INVITATION
>
> Some readers will be throwing their hands up: how naive — business runs on efficiency, of course you reinvest the saved time. I want to leave that here rather than resolve it, but let me put my own position on the table. Paying knowledge workers and managers by time rather than by output was always a concession to practicality, never a principle. These new technologies are a good reason to rethink that system rather than defend it by reflex. I'd genuinely like to know where you land — write to me.

So how do you build incentives that work? Two parts. **First, reward the behaviour, not just the output.** Stop taxing efficiency — set the explicit norm above, or people will rationally hide the gain. Reward multipliers rather than leaderboards: honour the person whose method others adopt, which is a team signal, not an individual contest. Make sharing cheap in the way just described. And legitimise experimentation by giving it a name and a slot (3M's fabled 15% time produced the Post-it; Google's 20% produced AdSense). **Second, make it count** — put it into the sheet you actually measure: a line for designing workflows that combine human judgement and AI, AI proficiency as a scored criterion in reviews and promotion, and explicit credit for reusable methods colleagues adopt. One concrete form: a visible contribution board that names the people whose help powered the most successful projects. Alongside all of this sits Amy Edmondson's distinction between praiseworthy exploration and blameworthy deviance — punish the first as if it were the second and you punish exactly the behaviour you are trying to embed.[6]

What I've seen in practice

The clearest case I've seen was not a training failure at all — and what made it instructive was how counter-intuitive it looked at first.

Picture a routine most organisations have some version of: every morning, an analyst builds a briefing by hand — pulling figures from five different systems, assembling them into a standard report, writing a short interpretation of what the numbers mean, and emailing it to a rotating list of stakeholders. Solid, valuable, and entirely manual.

In the pilot, AI took over the mechanical part. It pulled the five sources, drafted the report, and wrote the summary email. The human job shrank to the part that actually needs a human: a validity check — ideally in dialogue with the AI, asking why a figure had moved, whether the logic held, probing the edge cases — and then approving the send. It was faster, easier, and frankly more interesting work. Trained properly, genuine early wins, the numbers were there.

And then — the surprising part — several people quietly went back to building it by hand. The new way was demonstrably less effort, and they abandoned it anyway. On paper that makes no sense. It only started to make sense after a few honest conversations, which surfaced three reasons that had nothing to do with the tool:

- **The reward system still measured the old work.** The review sheet literally listed "independent creation and distribution of reporting" as a criterion. So the people who leaned hardest on the AI were, on paper, doing less of what they were formally rewarded for. Rational response: keep doing it by hand.
- **The "saved time" was taxed, not banked.** The team lead — reasonably enough — filled the freed-up 60–90 minutes a day with a backlog of manual master-data clean-up from a legacy system: monotonous, lower-status, and precisely not what these people wanted. Instead, using the AI came to mean earning yourself a worse task.
- **They lost the social contact — and feared for their standing.** Previously, after each send, stakeholders would call to discuss the numbers — and those calls were where the analyst was visibly the expert. Now those conversations happened between the stakeholders and the AI. Quiet fear of reputational loss, dressed up as "the old way worked fine."

Notice that none of the three is about the technology, and none of them shows up unless you ask. The fix, once we understood them, was three matching moves:

- **We rewrote what the review rewarded.** Out went "creation and distribution of reporting"; in came the contribution the new process actually needs — making the AI-supported output trustworthy: building the checks that catch the errors people used to catch by eye, and defining what "correct" means for the system. Using the AI stopped costing you credit and started earning it.
- **We let them bank the saved time in work they valued.** The freed-up hour went to the process-improvement projects they actually wanted — and those were named and recognised in the review, so the reclaimed time was a reward, not a punishment.

- **We gave the lost status a new, visible home.** Instead of fielding scattered post-send calls, the analysts ran a short weekly session walking stakeholders through how to interrogate the AI's output themselves — turning them from the person who produced the answer into the person who taught the room to trust it. Same expert standing, more leverage, and social contact restored on higher ground.

We also retired the old template so the AI-drafted version became the default file, and anchored the validity check to the start of the morning so it rode on an existing routine. But those were the easy structural moves. The three that mattered were the human ones — and we'd never have found them from a dashboard. Once all of them pointed the same way, the behaviour held, because going back was finally harder, and less rewarding, than going forward.

Make the new way the easy way, reward it where people actually look — and go find the human reasons a dashboard will never show you.

What this means for you

Each layer of the organisation has a different job in making the new behaviour stick.



Three levels, three jobs — what each layer owns in making the new behaviour stick.

C-level builds the system. Redesign performance management to reward orchestrating AI, not just direct output. Set the explicit norm so efficiency is banked, not taxed. Put AI fluency into reviews, promotion and succession. And make your own AI use visible — you are the most-watched signal in the building.

Middle managers make it safe and visible. Make sharing cheap: show once, the team documents, the credit stays. Reward the multipliers whose workflows others adopt. Separate praiseworthy exploration from blameworthy deviance in how you react. And let early adopters tell their own story — peer proof travels further than a top-down message.

Employees anchor the routine. Attach the new tool to an existing trigger rather than leaving it beside the workflow. Expect two months, not two weeks — one missed day resets nothing. Share what works. And use the sanctioned experimentation time rather than treating it as indulgence.

One reaction underpins all of this and no sheet fully captures it: the thirty seconds after someone shows you their work. A sceptical question, visible disappointment, or genuine interest — that is one of the most decisive trust signals you send about whether hidden workflows come into the open. Get those thirty seconds wrong and no incentive design will save you.

The Hidden Workflow & the Incentive Behind It

30 minutes, on your own or with your team. Two parts, straight from the diagnostic I use in the room.

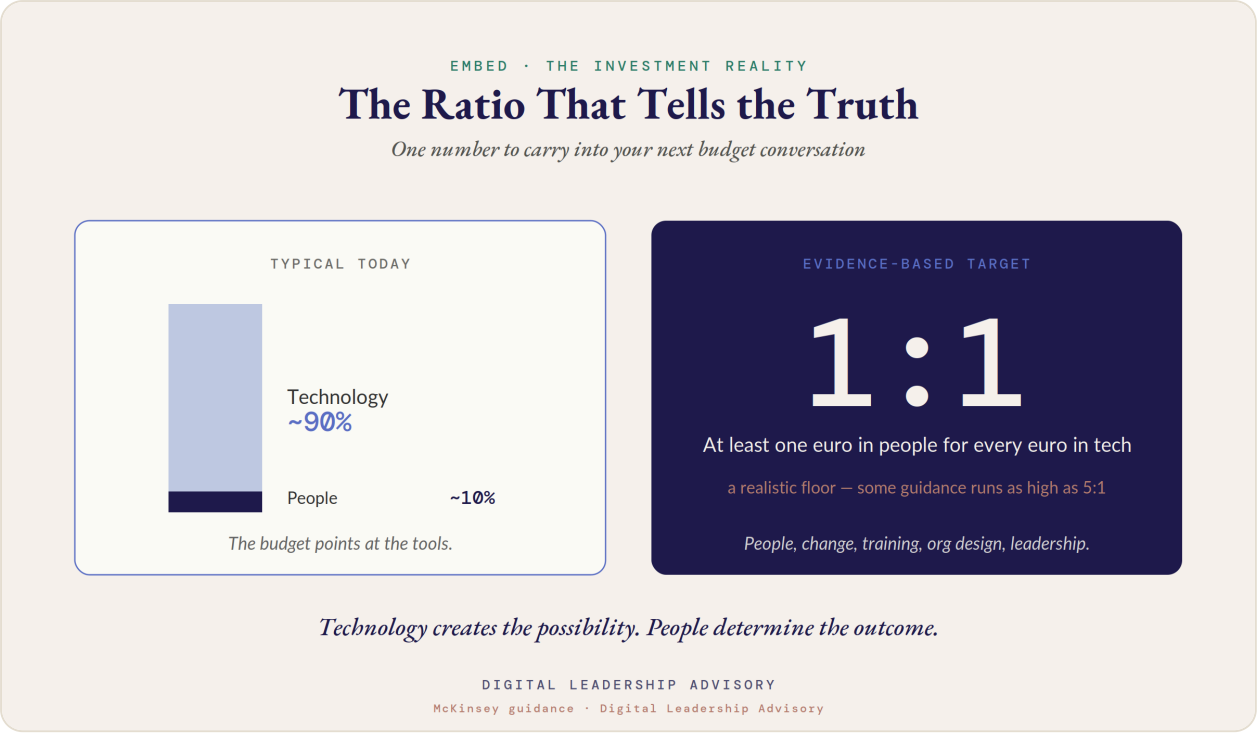
Part 1 — Hidden Workflows. Are there structures or unwritten rules in your organisation that could lead people to keep their AI workflows to themselves rather than share them? Be honest — and name one concrete thing you could do about it. (If you genuinely can't think of a single hidden workflow, ask why nobody feels safe enough to have a visible one.)

Part 2 — Incentive Design. What would a good incentive structure in your area look like — one that actually rewards sharing and AI-enabled value creation, not the old behaviour? Which of the four drivers above — reputational, workload, replaceability, support burden — is most alive on your team, and what is one change you could start next week to answer it?

And the part no structure captures: the next time someone shows you their work, what will your first reaction be in the first thirty seconds — and is it the one that earns you the next disclosure?

Closing the series

Four essays, one argument. Companies stall not because the technology underdelivers but because the organisation is not yet ready or able to do what the technology makes possible — the paradox we started with. Engage turns a vision into a direction people can act on. Enable gives them the ability to act. Embed makes the new behaviour the default, so it outlives the pilot.



The ratio that tells the truth — one number to carry into your next budget conversation.

Underneath all three sits one number worth carrying into your next budget conversation. Most organisations spend the overwhelming majority of their AI budget on technology and only a sliver on people — adoption, change management, training, org design, leadership. The evidence points the other way. Some guidance runs as high as five to one in favour of people;[4] my own rule from practice is more conservative — treat 1:1 as the realistic target, at least one euro in people for every euro in technology, a level I've also seen cited at McKinsey. Wherever you land, the ratio is the tell: it shows whether you believe the closing line of this whole series. Technology creates the possibility. People determine the outcome.

BEFORE YOU GO

This was the first full series from Digital Leadership Advisory — four essays on the human side of AI transformation. I hope it was useful. I'd genuinely like to know what landed, what you'd push back on, and what you'd like me to write next. Leave a comment, or reply — I read everything.

Sources

1 Lally, P., van Jaarsveld, C. H. M., Potts, H. W. W., & Wardle, J. (2010). How are habits formed: Modelling habit formation in the real world. *European Journal of Social Psychology, 40*(6), 998–1009.

2 Milkman, K. (2021). *How to Change: The Science of Getting from Where You Are to Where You Want to Be.* Portfolio.

3 Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025.* The University of Melbourne & KPMG.

4 Lamarre, E., Smaje, K., Levin, R., Singla, A., & Sukharevsky, A. (2026). *Rewired: The McKinsey Guide to Outcompeting in the Age of Digital and AI* (2nd ed.). Wiley. (Investment-ratio and operating-model guidance; a 1:1 people-to-technology floor reflects the author's own practice, also seen cited at McKinsey.)

5 Singh, B., Murphy, A., Maher, C., & Smith, A. E. (2024). Time to form a habit: A systematic review and meta-analysis of health behaviour habit formation and its determinants. *Healthcare, 12*(23), 2488.

6 Edmondson, A. C. (2011). Strategies for learning from failure. *Harvard Business Review, 89*(4), 48–55.