

RESEARCH WHITEPAPER

The Other Harness: What Turns AI's Possibilities Into Value

A sociotechnical proposal for judging, connecting and using AI-assisted work

Zartis Innovation Team

October 2026

Executive summary

AI has expanded what we can make and try. It has made it easier to produce an analysis, a piece of software, a proposal or a new service concept, often before an organisation has decided who might need it. Some of that work helps immediately. Some remains unseen, and some was never worth using. The question that now concerns us is no longer only how to build with AI. It is where the value appears, and what happens to the people who must recognise it.

Our Head of AI and Innovation, Adrián Sánchez de la Sierra, created Delegation Grade Architecture (DGA) as a research programme at Zartis. It has led us through three questions: can we delegate safely, can we delegate judgment, and does delegated work become value? Forge explores the controls around an agent. Deputy asks what it means to model a person's judgment. The third question leads beyond the agent to what we call *the other harness*: roles, approvals, relationships and attention arranged around the people who make and use the work.

This paper proposes an *absorption layer* for that third question: a versioned route from candidate work or an unmet need, through permission to discover a possible fit, to a human introduction, an authorised trial and observation of its effects. The route crosses three connected models with different sources of authority: a private individual, an employee acting within a role, and an organisation with needs and capabilities. It records what changed for the beneficiary, who carried the cost and what the people involved can still do next time. CommonField, a consent-based connection idea in the DGA programme, illustrates one part of the proposal.

The first implementation can be small. Take one recurring handoff. Ask what need it exists to serve, which judgment it protects, whose attention it consumes and what evidence would show that a change helped. Then ask a harder question: which human capacities would the new arrangement strengthen, and which are we willing to stop practising? A technical advance can enlarge tomorrow's possibilities only if people remain able to see and judge them.

1. The questions changed

In conversations led by our Head of AI and Innovation, questions from companies have changed faster than their organisations have. In 2023, many wanted to know whether they were late to AI. In 2024, discussion turned to how to build with it and who should build. Reliability became more pressing in 2024 and 2025; adoption followed. By 2025 and 2026, the question was increasingly difficult to answer with a demonstration: where is the value? This is a record of those conversations, not a survey of the market. Each new question appears when progress on the previous one reveals another limit.

The contrast in [McKinsey's 2025 survey](#) gives that question force. Eighty-eight per cent of respondents said their organisations regularly used AI in at least one business function, up from 78 per cent in the previous survey. Thirty-nine per cent attributed *some* company-wide impact on earnings before interest and taxes (EBIT) to AI, most often below five per cent. The survey had 1,993 respondents in 105 countries. Use and attributed EBIT are different reported measures, so the distance between those percentages is a puzzle rather than a calculated stock of unrealised value. It may contain delayed returns, customer gains a firm cannot capture, weak projects, or work that never reaches anyone who can use it.

Robert Solow's [1987 observation](#) about computers appearing everywhere except in the productivity statistics offers a useful warning against drawing an economic conclusion from technical visibility.

The analogy tells us to look beyond deployment. It cannot tell us which of today's explanations is right. To do that, we need to follow particular work and the people around it.

We began with a simple expression: $\text{value} = \text{potential} \times \text{absorption}$. The multiplication sign suggests a dependency. More things we *could* do will matter only if someone can turn a credible possibility into a useful outcome. It was a productive first question, but its first term needed repair. AI clearly expands what we can make and try; it has not been shown to multiply the number of opportunities that will pass a particular need and quality test. A generated output is a *candidate*. It becomes an *opportunity* when a beneficiary can name a plausible need and someone has reason to trust the fit. It becomes an observed result only after use and a comparison with the costs of making that use possible. An experiment undertaken to learn may have served its maker well without ever becoming a company project.

That distinction changes how we read the formula. Potential is discovered and sometimes created through conversation and redesign. Absorption is a set of actions and capacities, not a coefficient left over after value is measured. The actions include finding a need, judging quality, gaining permission, trying something and learning from what follows. The capacities include being able to notice an error or change the question. Judging fit can itself create an opportunity, so potential and absorption are not independent stocks we can simply multiply. If we define absorption as realised value divided by potential, the formula becomes true by construction and tells us nothing about where to intervene. The route through decisions replaces that literal calculation: we can ask which step failed and whether fixing it is worth its cost.

The three DGA questions make this development visible. First, an agent must be safe enough for its work to be considered. Then someone must decide what counts as a good result and which part of that judgment can be delegated. Finally, the work must find a purpose and a path into use. These are not stages every firm must traverse in order. A support team can gain from a useful assistant within an existing workflow. The questions describe where the scarce work may move as making an output gets cheaper: from production to verification, judgment and connection. The technical harness guides an agent through instructions and checks. The other harness gives a person a role, permission to decide, colleagues to reach and time to notice a need. Organisations design it, even when they never give it that name.

2. Can we delegate safely?

The first DGA question was about trust and control. When an agent acts through several steps, failure can take more forms than a wrong sentence. It can violate a constraint, announce completion before the work is done, proceed from a false premise, optimise a proxy at the task's expense, or degrade during a long session. In Forge we use five names for these modes: *constraint violation*, *false completion*, *false premise*, *reward hacking* and *long-session degradation*. They let us inspect a task and ask what evidence would expose each failure. The DGA programme also includes a trust-score model assembled from 147 signals. The underlying matrix is unavailable for independent assessment, so the count describes the scope of that internal account, not a validated accuracy rate.

Forge asks what an agent needs around the model in order to be trusted with more work. An obligation the agent cannot quietly rewrite differs from a preference in the next prompt. An explicit strategy differs from an invisible sequence of improvised steps. Evidence of a completed task gives a reviewer something to inspect beyond the agent's own assertion. These controls can make failures easier to find, but choosing

obligations, maintaining tests and judging the evidence all take time. As production accelerates, the ability to verify plausible action can become the scarce resource.

Consider an agent preparing a recommendation from a collection of documents. Requiring a citation for every claim may help with one error while leaving the question's premise wrong, a necessary document absent, or a recommendation aimed at the wrong outcome. A citation check asks whether the passage exists. A second check asks whether it supports the claim. A third asks whether the recommendation should be acted on. The checks can be composed differently according to the task's risk: a reversible internal draft may need sampling, while advice with a difficult-to-reverse consequence may need named approval. A useful harness exposes the contract of the task: what must remain true, what completion evidence counts, and who may approve an action.

The term *harness* has a technical life as well as the wider sense we use here. In a [2026 engineering account](#), Prithvi Rajasekaran describes planner, generator and evaluator roles for long-running application development, with handoffs that changed as models improved. A scaffold has to be judged alongside the model and task it serves. Forge applies that same design question to our research programme: which control is needed for which failure, and what new work does it create for the human reviewer?

Safety is an important achievement, yet a perfectly checked agent can complete a task nobody needed. This observation moves us from whether an output meets a specification to who chose the specification. A recommendation can satisfy its brief and miss a client's underlying problem. When that happens, the next boundary is judgment.

3. Can we delegate judgment?

Judgment appears when the criterion itself needs to be chosen or revised. A summary can be factually accurate while hiding the uncertainty on which a decision turns. A technical fix can pass tests while changing what its user valued. The human work in those examples includes noticing that the stated task and the real need have diverged. AI can assist with that work. The question is what the person continues to practise when the assistant begins to propose the criterion as well as the answer.

Deputy, our second line of enquiry, begins with a person's resolved cases rather than a generic job description. We can ask which distinctions that person actually used, which decisions can be delegated, and which must remain open to challenge. A past decision is evidence of someone's judgment, though it may also reflect missing information or a constraint imposed by a team, client or profession. The point of modelling it is to make the boundary discussable: what would a good deputy have to recognise, when should it ask, and who may correct the model of judgment? Those questions are more precise than saying that a role can or cannot be automated.

We recognise the cost of avoiding those questions in our own programme. Its founder, our Head of AI and Innovation, says his productivity has grown so much with AI that he has lost focus and that this affects the company. He believes he has probably produced more total value than he would have without AI, but we cannot measure that comparison. He puts the uncertainty plainly: "I do not know exactly what has changed, because I have not had time to stop and think about it." When an idea can become code almost immediately, we can postpone deciding whether it deserves a week of attention, a conversation, or no further work. That leaves us more material to judge and less space in which to form that judgment. Colleagues may experience some initiatives as useful; others may experience them as friction.

Consider a junior analyst who can now produce a convincing client answer in minutes. A senior colleague may catch today's false premise. If the junior never investigates why the client rejected an apparently sound answer, the team may lose a future source of insight. The immediate answer and the analyst's later ability to question one are different outcomes. An assistant can affect both, so we need to ask which part of the work the analyst still performs.

The mechanism we worry about is practice. We learn a distinction partly by having to make it. In a randomised study by Shen and Tamkin, 52 participants used an unfamiliar Python library; some learned with an AI assistant and others without. On an immediate comprehension quiz, the assisted group averaged 50 per cent and the unassisted group 67 per cent. This is a result for one learning task and a short horizon. It invites a design question: when the assistant solves the first difficulty, what part of the explanation does the learner still work through?

The answer can be different when teaching is part of the design. Kestin and colleagues tested a structured AI tutor in introductory physics. In the measured lessons, students learned more than under an active-learning classroom comparison and took less time. The assistant's prompts and sequence were designed to teach rather than merely hand over an answer. For a skill that matters later, we would assess a person's ability to act without the tool after an interval, alongside their speed with it now. That gives a team a choice about which practice it wants to preserve.

AI can also change what a group sees. In Doshi and Hauser's experiment, AI ideas improved the evaluation of individual eight-sentence stories, especially those by less creative writers, while assisted stories became more similar to one another. In a P&G field experiment published in 2026, AI-assisted individuals matched unaided two-person teams on early product proposals and generated ideas across commercial and R&D boundaries. The published main analysis contains 791 randomly assigned professionals. One study concerns short stories, the other early proposals. Together they sharpen the question of what a group will recognise and develop after each person has made a better object.

Tom Zahavy's 2026 ICML position paper, "LLMs can't jump", argues that current language models lack a particular *abductive leap*: proposing a new explanatory frame when familiar answers no longer fit. It is a position about discovery rather than a final measurement of model limits. In a different kind of evidence, the McKinsey HR Monitor 2026, based on a January survey of 1,303 HR professionals, recorded directional increases of 9, 6 and 3 percentage points in respondents selecting problem solving, creativity and reasoning among their five most important future skills. The survey records expectations, while Zahavy asks about a mechanism. Neither tells us how much judgment a particular company has lost. They help us name what to watch when work changes.

For that analyst, a structured tutor, a required explanation or a different split of cases could change the outcome. The components are small: a decision standard, a chance to challenge it, feedback from use, and protected practice for a skill someone expects to need again. A company can distribute these among people and tools in several ways. What matters is whether anyone retains the authority and ability to recognise the next unfamiliar case. The way we make work changes the people who must judge the next piece.

4. Does delegated work become value?

The third question begins with an object that looks useful and asks who can make use of it. A prototype may sit in a repository, a report may be read without changing a decision, and a tool may be tried once then abandoned. Each is an event; none by itself tells us what need was served. An output can fail because it is poor, because there is no timely need, because the people who need it never hear of it, or because using it would violate a legitimate boundary. A route through these decisions is what we mean by *absorption*. It can begin with a need that inspires work or with an experiment that later finds one.

In one quarter, the programme's founder produced roughly 150 codebases. About ten reached someone who used them. Much of this began from a personal motive: improving his agents or workflows, or learning something. He did not design every codebase for others, and we have not analysed why some travelled while others stopped. The ones that had a chance beyond their maker were those he took time to share, on Slack or in direct conversations. Sharing did not guarantee use. It was the path by which someone else could discover the work. The count is his account, without a common inventory or measure of net benefit; the lesson we draw concerns attention and connection, not a conversion percentage.


A preliminary Project NANDA report from 2025 hints at how much work can sit outside official organisational channels. In its selected sample of 52 organisations and 153 senior leaders, 40 per cent of companies had bought official LLM subscriptions, while workers in more than 90 per cent of the companies reportedly used personal AI tools for work. That second percentage counts *companies with such use*, not workers. Informal use can coexist with limited official provision, leaving a firm with work it may benefit from but struggle to see, govern or reuse.

The phrase *absorptive capacity* has a history. Cohen and Levinthal described a firm's ability to recognise, assimilate and apply valuable external knowledge, shaped by what it already knows. Zahra and George later distinguished potential and realised absorptive *capacity*. We extend their question to internally made, AI-assisted work: how does a firm become able to notice and use what its people can now create? The terms are related, though our "opportunity" concerns a particular possible use, whereas their "capacity" concerns the organisation's ability. This distinction matters when a team has many candidates but lacks the shared knowledge to judge them.

We explored that situation with a **synthetic world with invented parameters**: a consultancy of ten people or agents, clients, needs and limited attention. It reports 83,200 runs across 200 simulated worlds and 13 variants. In one comparison, works produced rose from 82 to 213 while mean needs served moved from 84.4 to 83.1. The model exposes a possibility: output and needs served can move in different directions when the rules governing contact and attention are held apart.

More production, no improvement in needs served

Synthetic world · invented parameters



83,200 runs · 200 worlds · 13 variants.

Different units and scales; model code and readout unavailable.

Figure 1. Synthetic world with invented parameters, comparing without AI and with AI. Works produced and mean needs served use separate units and scales.

If more work is made, the scarce resource can become finding a meaningful connection. Yet an indiscriminate feed asks every recipient to screen more material. The original formula now has to do more work than two factors can hold: creation changes what people know; connections cost attention; use can fail a later test; and a person may need practice to spot the next opportunity. Rather than defend a number called absorption, we can inspect a route: **candidate or need -> permitted projection -> proposed fit -> mutual introduction -> authorised trial -> observed effects and costs**. A person may see a possible use privately before projecting anything across models. Feedback may return to the work, the need or the person's skill. That route is the beginning of an engineering proposal.

5. Replace or reimagine?

Once a connection bottleneck is visible, we can ask whether the work around it should change. When a model makes an existing task faster, replacing part of that task can be an excellent first move. The larger design question appears when the surrounding workflow still reflects an older constraint. Suppose a specialist once had to prepare every first draft because producing a credible one took scarce expertise. If drafts become cheap, the specialist may now receive far more material to review. Removing review would expose users to error; preserving every old step could turn expertise into a queue. The choice is to discover what the review protects, then distribute checks, approval and learning in a new arrangement where the old form no longer serves its purpose.

Paul David's 1990 account of electricity and the factory makes this more concrete. An early electric motor could replace the power source while shafts and belts kept the factory's layout. Distributed unit drive later allowed machines, buildings and material flow to be redesigned. The gains were learned through reorganising the factory, not merely changing the motor. David also understood that such complementary changes take time and knowledge. Solow's productivity puzzle was therefore never answered by counting installed motors. The productivity J-curve offers a contemporary way to think about investment in processes and skills preceding measured gains; McKinsey Global Institute's 2026 discussion, especially Exhibit 2, uses electricity to illustrate unequal speeds of technical and organisational change.

Redesign does not follow automatically from the parallel. In a published study of 5,172 customer-support agents, AI assistance accompanied about 15 per cent more issues resolved per hour after a staggered rollout, with larger gains among less experienced workers. An email field experiment involving 66 firms

and 7,137 workers found about two fewer hours a week on email among tool users in the second half of the trial, without a detected change in task composition from individual-level provision. Those are useful local results. If the saved time has a valuable use and the duties around the task still fit, a firm may have little reason to rearrange everything around it.

The review queue also shows why a collection of task savings may be too narrow an ambition. Acemoglu's 2025 model projected at most a cumulative 0.66 per cent increase in total factor productivity (economy-wide efficiency after accounting for labour and capital) over ten years through the then-estimated task-exposure and saving channels. Acemoglu and Restrepo distinguish that kind of automation from the creation of new tasks. In the P&G experiment, a person assisted by AI produced early proposals across commercial and R&D boundaries, candidates for work a single role might previously have missed. Their later commercial value remains unknown.

For the specialist's queue, the choice is where to place judgment after drafts become cheap. Agrawal, Gans and Goldfarb's framing directs attention to judgment and action that complement cheaper prediction. Porter's account of fit among activities asks whether those choices form a coherent service for a customer. Moving routine checks before the queue, keeping approval for consequential uncertainty and giving rejected drafts back to their makers might create such a fit. It is a design to compare with the present queue, rather than an economic conclusion inherited from the studies.

Roles themselves are part of the existing harness. Smith's division of labour explains why specialisation can increase productive capacity; Simon's bounded rationality explains why people need simplified, distributed decisions. We use those as concepts for interpreting roles designed before present AI capabilities. A title may bundle execution, tacit quality control, knowledge of the customer, authority and a bridge to another team. If AI changes the first component, the others do not vanish. Some may become more important because the volume of plausible work has increased.

Connections are part of that design. Burt's 2004 study of structural holes associated brokerage between otherwise separated groups with access to different ideas and favourable evaluations in one large company. A broker can notice a use that neither group sees alone; an organisation-wide feed can instead overwhelm its readers. David also described the screening burden created when distributing information gets cheap. In the programme's **synthetic world with invented parameters**, reassigning a person met the model's success criterion in 0 of 200 simulated worlds; making work visible did so in 43 of 200. The success criterion was not available in the supplied model material, so these counts describe the synthetic comparison without establishing an operational outcome. The model also reports a 12 per cent increase in an attention measure whose definition and denominator were unavailable. The exercise sets up the design problem: reveal a useful connection without making everyone screen everything.

McKinsey's 2025 survey calls roughly six per cent of respondents *AI high performers*, defined by reported significant value and at least five per cent AI-attributable EBIT impact. They were nearly three times as likely as others to report fundamental workflow redesign. The association is a reason to investigate complementary changes, not a causal estimate of redesign. A leader can start more concretely: for a specific handoff, separate the effort a model has made cheaper from the right, relationship or skill that the handoff exists to protect. Then decide whether a new connection or review path improves the result enough to pay for itself.

6. The absorption layer

An *absorption layer* keeps a record of a route: what was made, what was needed, who agreed to connect them, and what happened. Deputies use that record to look for possible uses as work and needs change. They can propose a connection, bring evidence to a human decision and return what a trial taught. Imagine a developer who makes a small tool to compare approaches to a client problem. She made it to improve her own work. A different team has a related need; neither knows about the other. The layer should notice the possibility without requiring an announcement to everyone.

The original diagram gives this route **one person in three connected models**. The *individual model* belongs to the person. It can hold aspirations, strengths, a skill she chooses to keep practising and her choices about delegation; it is private by default. The *employee model* is governed by the person and the organisation together. It describes work within a role, a Deputy's model of judgment, handoffs, approvals and an attention budget. The *organisation model* holds needs, capabilities and claims about available work under organisational and client restrictions. These models share some identities, but they have different authority over what may be seen and done.




Figure 2. Reading map of the active layer. Plate A in Appendix A reproduces the original ontology in full. The map adds proposed Deputy operations to the original relationships.

Plate A calls an unconnected output *latent potential*: it may have a use beyond its maker before anyone can name one. To work on that possibility, the layer records a particular version as a *candidate*. A named need and a judged fit make one possible use explicit. Sometimes the need comes first and a Deputy looks for the right work or person. CommonField's graph connects these things with people and capabilities. The original arrows also link delegation to possible erosion and learning opportunities to practice. Plate A places Deputy in the employee model. We propose a *private Deputy*, which helps a person choose what to offer from her own model, and a *role Deputy*, which sees only descriptions released for a team purpose. The layer invokes five functions as authorised changes and later observations occur:

- **Match** runs when a candidate, need or relevant capability changes. It proposes a connection and shows the evidence that led to it.
- **Fit check** runs before an introduction. The people who know the work and the need inspect quality, timing and restrictions, and may reject the proposal.
- **Receipt** runs when an authorised trial ends. The recipient describes whether it helped; the team compares that account with the baseline, errors, attention and maintenance costs.
- **Erosion check** runs after a period of delegation. Where a protected skill matters, an unaided task or other suitable evidence is compared with a pre-delegation baseline to see whether the person or team can still make the distinction.
- **Learning offer** follows an erosion signal or a new growth opportunity. It proposes deliberate practice and later checks whether ability was retained.

The graph proposes; people decide whether to connect and use the work. The five functions turn a static map into a learning loop, while the person's choices about protected skill remain hers.

These functions need a small, versioned memory. A first pilot could use six kinds of record, perhaps in a form and event table before building a graph:

Record	What it must hold
Candidate	Work, version, maker, original purpose and use restrictions.
Need	Beneficiary, problem, quality test, version and owner.
Role and authority	Who can judge fit, agree to an introduction and authorise a trial.
Projection grant	Which description may cross a model boundary, for which purpose, audience and period.
Fit proposal	Candidate and need versions, reason for the link and the human decision.
Trial and observation	Scope, baseline, separate use approval, receipt, effects, costs and later capability check.

Both candidate and need require independent permission to enter cross-model discovery. An introduction needs both sides' agreement; a trial needs its own current right to use the work. When a permission changes, the layer stops new discovery from relying on it. The technical appendix gives the expiry, revocation and audit rules. These are boundaries on the active functions, not their purpose.

In the developer's case, her private Deputy can suggest offering a short description of the tool while keeping her personal reason for building it private. A role Deputy notices the team's need in the authorised graph and proposes a match between those versions. The developer and team examine the fit and may agree to talk. Their owners then decide whether a limited client trial is allowed. The recipient's receipt might say the tool helped them reject an unsuitable approach. Observation asks whether that benefit outweighed review and maintenance work. If repeated use also removes the analysts' practice of making the initial comparison, an erosion check can lead to an offer of unaided comparisons with feedback. The offer creates an opportunity to learn; later performance shows what was retained.

Parts of this design have antecedents. [Workday's documentation](#) describes Talent Marketplace matching, match analysis and feedback; its Career Hub also suggests learning content. [Glean's Knowledge Graph documentation](#) describes links among people, content and activity with item-level permissions. The proposed extension follows a versioned AI-assisted artefact or unmet need across person, role and organisation models, with independent permission to make each discoverable, human-approved use, a recipient's account of help and cost, and a delayed check on a skill delegation may erode. A team could compare Deputies watching authorised changes with a trusted human broker handling the same events, counting useful introductions, rejected matches, attention, outcomes and retained skill.

The route gives the formula a more useful role. It shows where potential found a need, where a connection failed and what a trial changed for the people involved. It also lets an organisation ask whether today's gain is leaving it able to recognise tomorrow's opportunity.

7. Where the hypothesis could fail

The strongest test of the absorption idea is whether it explains anything after simpler alternatives are counted. First, value may already be appearing somewhere the usual firm measure misses. Time saved might improve a customer's experience without increasing EBIT; complementary investment might precede the return reported in accounts. The [productivity J-curve](#) makes a delay plausible. Following a use case through customer outcome, worker time, firm cost and margin over an agreed period would show whether a supposed gap was partly a measurement problem. The layer should be judged against that fuller account rather than against a survey contrast.

Second, some output should remain unused. Some of those experiments might have fulfilled their purpose by teaching their maker something. Other artefacts may be incomplete, poor quality or irrelevant. Treating every unshared codebase as an opportunity would reward production for its own sake. A fair test would register candidates before knowing which will succeed, including those their makers choose not to share. It would record original purpose, a plausible beneficiary and a quality judgment where one can be made. Some exploratory work will reveal a need later, so the record should allow that route without retroactively declaring every unused object valuable.

Third, a new connection can impose more work than it returns. We may see a useful initiative where a colleague sees interruption, duplicated effort or another maintenance obligation. The synthetic model showed that trade-off under its invented rules. A field comparison should count the sender, reviewer, recipient and affected colleagues, and should compare a specific implementation with both usual discovery and a simple targeted introduction. It should record rejected matches and abandoned trials as carefully as successful ones. Otherwise selection by enthusiastic makers and recipients will make any matching system look better than it is.

Fourth, useful value can arise without a new organisational layer. The support and email studies show local gains within largely existing patterns of work. A team with a clear task and beneficiary may need little beyond a good assistant and its normal controls. A graph that adds no better fit, no clearer authority and no net benefit after its costs would be the wrong answer for that team. The proposed architecture also has to survive its own governance tests: can a person correct or withdraw a claim, can an employer see a private aspiration indirectly, can a client restriction be enforced, and can a recipient refuse without consequence? A design that fails those questions would make the other harness worse.

Finally, the human-capacity claim can fail in a useful way. A skill may be rationally ceded when the tool performs it reliably and the released attention develops something more valuable. Structured tutoring can

improve learning, and technical controls may carry some checks more consistently than an old role did. We should identify the *particular* capacity a future decision will require, test it after a delay without assistance where appropriate, and watch how it moves across a team. If necessary judgment remains strong while delegation grows, the erosion part of our account must narrow. What matters is the organisation's ability to recognise and challenge a plausible but wrong answer, including through distributed expertise rather than one heroic individual.

A prospective study could give a human broker and the Deputies the same eligible changes. In one arm the broker proposes connections; in the other the Deputies do, with both passing through the same human fit, introduction and trial decisions. Compare useful introductions, beneficiary outcomes, attention costs and retained skill.

The formula with which we began has changed. Value = potential × absorption distinguished making something from making a difference. The route asks: what need might the work serve, who may judge and authorise it, what happened after use, and what can the people involved do next? The other harness is the arrangement in which those answers are found. We can redesign it piece by piece once we know what each part protects. The final decision is which capacities we are willing to cede, and which we want our organisations still to have when the next question arrives.

Source and method note

Our account of DGA, Forge, Deputy and CommonField draws on the Zartis research programme led by Adrián Sánchez de la Sierra. The changing client questions and codebase count are his observations; the consultancy simulation is a synthetic model developed within the programme. The absorption-layer image presents a design proposal from the programme; its privacy and business effects have no field evaluation here. The simulation is a synthetic world with invented parameters: it reports 83,200 runs across 200 worlds and 13 variants, but the model 1.2.1 readout, code, attention denominator, success criterion and sensitivity analysis were unavailable. The 147-signal matrix was also unavailable. The NANDA report is preliminary; its item denominators are unavailable and the 90 per cent label varies between prose and exhibit. The published P&G analysis uses 791 participants; 776 is the earlier working-paper and complete-survey count. HR Monitor changed its country coverage between waves, which makes the stated shifts directional. Public empirical, economic and conceptual sources are linked at the claims they support. Where a cited study concerns a bounded task, its stated outcome is used for that task rather than as a company-wide effect.

References

- Acemoglu, D. (2025). "The simple macroeconomics of AI." *Economic Policy*.
- Acemoglu, D., and Restrepo, P. (2019). "Automation and new tasks." *Journal of Economic Perspectives*.
- Agrawal, A., Gans, J., and Goldfarb, A. (2016). "The simple economics of machine intelligence." *Harvard Business Review*.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025). "Generative AI at Work." *Quarterly Journal of Economics*.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2021). "The Productivity J-Curve." *American Economic Journal: Macroeconomics*.
- Burt, R. (2004). "Structural holes and good ideas." *American Journal of Sociology*.
- Challapally, A., et al. (2025). *The GenAI Divide: State of AI in Business 2025*. Project NANDA. Preliminary report; accessible copy of original text.

- Cohen, W., and Levinthal, D. (1990). “Absorptive Capacity.” *Administrative Science Quarterly*.
- David, P. (1990). “The Dynamo and the Computer.” *American Economic Review*.
- Dell’Acqua, F., et al. (2026). “The Cybernetic Teammate.” *Organization Science*.
- Dillon, E., Jaffe, S., Immorlica, N., and Stanton, C. (2025). “Shifting work patterns with generative AI.” NBER Working Paper.
- Doshi, A., and Hauser, O. (2024). “Generative AI enhances individual creativity but reduces the collective diversity of novel content.” *Science Advances*.
- Glean (documentation accessed October 2026). “Knowledge Graph.”.
- Kestin, G., et al. (2025). “AI tutoring outperforms in-class active learning.” *Scientific Reports*.
- McKinsey & Company (2025). “The State of AI in 2025.”.
- McKinsey & Company (2026). *HR Monitor 2026*.
- McKinsey Global Institute (2026). *The AI economy*.
- Porter, M. (1996). “What is strategy?” *Harvard Business Review*.
- Rajasekaran, P. (2026). “Harness design for long-running application development.” Anthropic Engineering.
- Shen, J., and Tamkin, A. (2026). “How AI impacts skill formation.” Anthropic.
- Simon, H. (1955). “A Behavioral Model of Rational Choice.” *Quarterly Journal of Economics*.
- Smith, A. (1776). *The Wealth of Nations*, Book I, Chapter I.
- Solow, R. (1987). “We’d better watch out.” *New York Times Book Review*.
- Workday (documentation accessed October 2026). “Using the Talent Data and Career Development.”.
- Zahavy, T. (2026). “Position: LLMs can’t jump.” ICML position track.
- Zahra, S., and George, G. (2002). “Absorptive capacity: A review, reconceptualization, and extension.” *Academy of Management Review*.

Appendix A. Original absorption-layer ontology

Plate A keeps the original relationships and labels. Read across three authority contexts: a private individual, an employee with a role, and an organisation with needs and capabilities. Role, task and output can meet a need through a graph-proposed connection and two-sided agreement. Green paths offer practice; red paths mark possible personal and collective erosion. These are proposed relationships.

Plate A calls unconnected output “latent potential”: it may have a use beyond its maker. The implementation records a particular version of that output as a candidate. A named need and judged fit make one possible use explicit. The recipient’s confirmation records experienced help; an outcome check adds effects and costs when assessing value. An offered learning opportunity creates a chance to practise, while later evidence shows what was retained.



Appendix B. Governance rules

Cross-model discovery works only on descriptions released for a defined purpose; a private Deputy may help its person decide what to offer before any release. A candidate and a need each require their own projection grant when they cross an authority boundary. The grant identifies the source version, permitted fields, destination, audience and expiry; source owners and relevant role or client rights holders clear what they control. A shared identity between the private, employee and organisation models grants no access by itself. A Deputy can inspect authorised changes and suggest a fit, but it cannot grant a projection, introduce people without their agreement, authorise client use or certify an outcome.

Fit and introduction records retain the exact candidate and need versions, the discovery grants that allowed the proposal, and the authority and rule versions applied. These references are provenance, not continuing use rights. A trial records separate current approval for its scope, baseline and stop rule. If a discovery grant expires or is withdrawn, new matching and pending introductions stop. The layer cancels queued notices and makes copies in its controlled indexes and caches inaccessible. An authorised trial continues only while its separate use rights remain valid; a revoked use right suspends it until reauthorisation.

Changing a source, grant, authority decision or rule makes dependent proposals stale. Refusals remain restricted; aggregate patterns may improve matching without becoming personnel scores. Minimal audit and human exposure records remain because a message already seen cannot be withdrawn. A pilot must test whether match explanations or notification timing reveal private aspirations. People may correct or withdraw their private offers; client-derived work remains under client authority.

Evaluation contract

Both comparison arms receive the same eligible changes, source, client and person permissions, refusal rules, and human fit and use decisions. The trial records missed and false matches, recipient outcomes, attention, maintenance and later capability, including when an introduction is refused or a trial abandoned. It should attempt to withdraw a discovery grant, revise a need and explain a proposed match to someone who knows the employee. A test that introduces the record and permission package together with Deputy proposals measures that combination; a later trial could isolate the contribution of automation. A separate assisted-versus-unassisted comparison under the same prospective need and quality criteria would test whether AI expands qualified opportunities rather than merely attempts. A null or negative result would limit the case for this implementation in the setting tested.