

PAPER


Intelligence Pays

How leaders turn cheaper intelligence into
better work and lasting value

Colin Price and Sharon Towe

October 2026

Part of our series on leadership and organisational performance



Imagine an executive team reviewing its AI programme. Licences have been issued, the training was well attended and several colleagues can demonstrate something impressive. Then the finance director asks where the improvement appears in the business. The customer is still waiting three days for a quotation. The same people are checking the same information. Nobody has stopped producing the weekly report.

The technology has moved faster than the organisation.

This imagined scene captures a real management distinction. AI can make an activity quicker while leaving the business scarcely changed. A better draft still waits for approval. A faster answer creates more work for the colleague who must verify it. Reports become longer because they are easier to produce. Unfortunately, the people expected to read them have not received the same upgrade.

Our argument is straightforward. Choose the work you want to improve, then change the conditions in which people do it. That means dealing with information, decisions and economics, but also with workload, professional pride and the understandable fear of becoming less useful. These are part of the work of adopting AI. Leaving them unresolved puts the return at risk.

Published studies and company accounts show both substantial opportunities and disappointing results. The difference cannot be explained by enthusiasm alone. Leaders need a credible bargain with employees, practical ways to learn and someone accountable for turning improvement into value. They also need to examine how their own behaviour helps or obstructs that change.

The organisation determines what technology becomes

In *Accelerating Performance*, we developed META: Mobilise, Execute and Transform with Agility. Applied here, it asks whether the organisation can establish a worthwhile purpose, deliver reliably, change its capabilities and adjust as it learns. The common thread is time to value. Buying a faster tool helps only if the organisation can put it to useful work.

Beyond Performance, which Colin co-authored with Scott Keller, examines organisational health: the ability to align, execute and renew in ways that sustain performance. An impressive pilot demonstrates a possibility. Reproducing the benefit, discovering its limits and improving the method require an organisational capability. That is the more durable prize.

Neither book provides a percentage return to put into an AI business case. We are applying its management logic to a new problem. Our proposition is that useful practice is more likely to endure when the technology, the work and the conditions people experience support one another. The result still depends on the application and must be tested.

This also makes the separation between technical change and human change rather artificial. Who can see the information? Who may act on the answer? Who feels able to say it is wrong? A system embodies decisions about all three. Technology and organisation are already intertwined by the time anyone attends the launch.

Three conversions determine realised value

Realised value = Adoption × Improvement × Capture This is a conceptual formula, not a calculation of financial return. Adoption converts capability into repeated use; improvement

converts use into better work; capture converts better work into realised value. Strength in one factor cannot compensate for neglect of another.

The first is from capability to repeated use. An employee needs approved access, relevant information, time to practise and a reason to persist when the first attempt disappoints. A licence provides a starting point. Training attendance tells you little about what happens when a customer is waiting and the familiar method feels safer.

The second is from use to better work. Does the new method improve accuracy, quality, effort, elapsed time or the customer's experience? Include the colleagues who receive and check the output. A team can use AI frequently while making the whole process slower. A rising usage chart would make that problem harder to see.

The third is from better work to realised value. Released capacity might support more customers, clear a backlog, avoid recruitment or remove spending. Someone must decide which, and make the necessary changes to demand, staffing, service or pricing. Multiplying minutes saved by salary cost does not do that work.

The remedy depends on where the conversion breaks. Missing access needs an operational decision. Unreliable output needs a better method or a different application. A genuine improvement that never reaches the customer or the accounts needs a business owner prepared to act. Another communication campaign will not resolve all three.

“

The technology has moved faster than the organisation.

Begin the executive review there. Establish what is preventing a worthwhile change, who can remove the obstruction and what evidence would show that they have. Adoption becomes useful when it changes the work.

The AI value continuum

Six positions connect wider access to intelligence with new ways to serve customers.

1 Universal access	2 Embedded assistance	3 Workflow redesign	4 Operating model redesign	5 Enterprise replatforming	6 Business invention
Give people approved access to general purpose AI.	Build AI into recurring tasks and everyday systems.	Redesign a complete process; people manage exceptions.	Change linked processes, roles, teams and decision rights.	Rebuild an existing company's systems, operations and offer around AI.	Create a new business or a new way to satisfy demand.
VALUE SOUGHT					
Better individual work; dispersed time savings.	Repeatable task gains; capacity for growth.	Fewer steps, queues and handovers.	Less structural work and duplication.	A different cost base and customer proposition.	New customers, unmet needs and viable new revenue.
RELATIVE INVESTMENT AND CHANGE EFFORT					
Lower	Low to moderate	Moderate	High	Very high	Varies; demand and business model need testing
WORKFORCE AND ECONOMIC IMPLICATION					
Usually little direct headcount reduction.	Roles combine only if sufficient work disappears.	Targeted staffing reductions become credible.	Material scope in affected operations.	Substantial scope across the existing company.	Creating value and new work; staffing reduction is not the main test.

From improving existing work to creating a new business. Positions describe different ambitions, not a compulsory sequence. Investment and workforce effects depend on the application. Source: Authors' synthesis.

Choose the change the business needs

The AI value continuum distinguishes six investment choices. At one end, people gain access to a useful tool. Further along, the company rebuilds how it operates. At the far right, it creates a new business or a new way to satisfy customer demand. Confusing those ambitions creates unrealistic economics and an employee promise that the organisation cannot keep.

Across the first five positions, the depth of work redesign generally increases, bringing greater integration, disruption and scope to remove complete blocks of work. The sixth changes the question: what could we offer that was previously impractical or uneconomic? It introduces

uncertainty about demand and the business model, and need not involve more legacy investment. This is a conceptual framework, not a forecast or a timetable. Different parts of one business may belong at different positions.

1 Universal access makes intelligence widely available

Approved access to leading general-purpose models can help employees draft, analyse, summarise and explore ideas. The fall in usage prices is striking. Stanford's 2025 AI Index reported that the cost of querying a model matching GPT-3.5 on a specified benchmark fell from \$20 to \$0.07 per million tokens between November 2022 and October 2024. That is a model usage price, not the cost of changing a business.

The benefit may be better work, less overload or slower recruitment. Time is often released in fragments across different people and days. Those fragments do not automatically combine into jobs that can be removed. Large staffing reductions are a weak starting assumption for a licence rollout.

Human challenge: Fear of change can keep people with the familiar method, while senior leaders hand responsibility to IT. Leaders need to use the tools themselves, make time for practice and own what changes in the work. Access is a technical provision; adoption requires a credible reason to try.

2 Embedded assistance changes recurring tasks

Here AI becomes part of a role's everyday systems: preparing a proposal, retrieving approved information or completing a meeting record. Morgan Stanley's Debrief, discussed later, makes this tangible. The benefit can repeat because the assistance sits inside recurring work.

The organisation must invest in information, system connections and proportionate checks. Capacity may support growth or avoid hiring. Combining roles becomes possible only where enough work has disappeared, in the right places, to change how the business is staffed.

Human challenge: Colleagues must change familiar habits without surrendering professional judgement. Some will check every word until the saving disappears; others will trust a fluent answer too readily. Practise the task, agree proportionate checks and make responsibility for the result clear.

3 Workflow redesign removes whole blocks of work

At this point, leaders redesign a complete process. AI handles eligible routine cases; people manage exceptions and consequential judgements. Queues, handovers and unnecessary steps can disappear. Staffing reductions become more credible because the requirement for human effort changes across a defined operation.

Klarna reports that its assistant handled 80% of customer-service chats in 2025 and delivered approximately \$59 million in cost savings. It describes the workload as equivalent to more than 850 full-time agents, while retaining human support. These are company-reported figures. Workload equivalence does not establish that 850 employees left. The case illustrates process economics, rather than the rebuilding of an entire enterprise.

Human challenge: Removing steps can also remove work people value, expose colleagues to redundancy and unsettle relationships between teams. Involve those who know the process, explain which decisions remain open and provide credible support for changing roles. People need to understand what replaces the work they are being asked to give up.

4 Operating model redesign changes roles and boundaries

Several connected processes now change together, along with responsibilities, team sizes, management layers and shared services. The workforce may now combine people with AI agents that can carry out tasks and take actions within agreed limits. A services business might connect enquiry, quotation, scheduling and billing through this shared arrangement, with people accountable for commercial commitments.

This creates greater scope to remove structural work, alongside new demands for expertise and control. It requires integration, changes to roles, retraining and the retirement of old arrangements. Automating each department's contribution separately leaves many of the costs between departments intact.

Human challenge: Leaders must learn to lead a workforce of people and AI agents. Decide which work to delegate to agents, where human judgement is essential and who can intervene when an agent is wrong. Colleagues need to learn how to direct, check and challenge agents without checking everything twice or surrendering responsibility. Managers may also lose teams, authority or status. Recognise them for the performance of the whole service, and keep accountability for outcomes with a named person.

5 Enterprise replatforming rebuilds the company on AI

Here AI becomes a foundation for an existing company's operations and customer proposition. Data, applications, processes and roles are redesigned together. Major legacy systems may be replaced. Routine execution and coordination work may disappear across substantial parts of the business.

The potential is considerable, as are the transition costs and execution risks. Data migration, parallel running, workforce changes and new capabilities all need funding and attention. Growth may absorb released capacity; some roles may disappear and others emerge. Leaders owe employees an honest account of those possibilities, with a clear distinction between decisions made and questions still open.

Human challenge: People have to keep today's business running while learning a different one. Fatigue, uncertainty and attachment to hard-won expertise can undermine the transition. Sequence the change, stop competing demands and be candid about future roles. Leaders must also relinquish arrangements that once made them successful.

6 Business invention creates new ways to meet demand

AI can make a service possible at a price, level of personalisation or scale that previously ruled it out. The opportunity may be a new business, a different commercial model or a way of reaching customers who were never served. The underlying need may already exist; what changes is the customer's ability and willingness to act on it.

Imagine an AI-enabled service that gives a small business continuous operational advice and helps carry out agreed actions for an affordable subscription. That is an illustration, not an established result. The proposition would need to prove that customers use it, trust it and pay enough to cover delivery and support. Its value would come from creating a viable market, rather than removing headcount from an existing process.

Human challenge: Embrace paradox to reframe the opportunity. An apparently unavoidable choice between affordable scale and personal service may reveal a new business model. As we argue in *Accelerating Performance*, look for a design in which the two objectives reinforce one another. AI

might make routine support inexpensive enough to fund expert attention when it matters, allowing a service to reach customers previously excluded by its cost. Leaders must help teams challenge the assumed trade-off, then test the demand, capabilities, investment and economics needed to make the new proposition work.

Make deliberate choices across the portfolio

The continuum is not a ladder every company must climb. Choose the change capable of producing the result you need. A useful assistant may be entirely sufficient in one area, while a new service requires a different business model elsewhere. Ambition should determine the investment, authority and transition support required. For business invention, the test includes demand, repeat use and viable economics, as well as reliable delivery.

The workforce conversation must match that choice. Promising that every role is safe while budgeting for structural reductions is an obvious contradiction. So is presenting modest task assistance as enterprise transformation. People can work with difficult truths. They have less reason to commit when the organisation gives them incompatible ones.

The evidence makes the opportunity more specific

Useful evidence tells us where gains arise, who benefits and what can go wrong. Controlled studies help separate an effect from enthusiasm. Company accounts show what organisations report at scale. Neither supplies a universal productivity number. Read each case for the decision it helps you make.

Customer service shows how expertise can spread

In *Generative AI At Work*, published in 2025, Erik Brynjolfsson, Danielle Li and Lindsey Raymond studied the staggered introduction of an assistant among 5,172 customer-support agents at a business-software company. It suggested responses and information during customer conversations; agents could edit or ignore them. Access increased issues resolved per hour by 15% on average, with the largest gains among less experienced and lower-skilled colleagues.

The most experienced and highest-skilled agents saw small speed gains and small quality declines. The researchers also found evidence of learning that persisted when the tool was unavailable. This was one company and workflow, studied through a staggered rollout rather than fully randomised deployment. Its findings belong to that setting.

The management implication is worth pursuing. Expertise is more valuable when people can use it at the moment of need. Making a good answer accessible during a customer conversation changes its usefulness. Publishing another manual may leave the distance between knowing and doing much as it was.

But the experts whose judgement informs the method need a meaningful future within it. Ask them to define quality, examine difficult exceptions and develop others. Recognise the performance they enable across the organisation. Otherwise the invitation to share expertise can sound remarkably like an invitation to surrender their standing.

Procter and Gamble illustrates a change in collaboration

The *Cybernetic Teammate*, published in *Organization Science* in 2026, reports a field experiment involving 791 professionals at Procter & Gamble. Fabrizio Dell'Acqua and colleagues compared individuals and cross-functional pairs tackling product-innovation challenges, with and without

generative AI. Individuals using AI produced work comparable in quality to unaided two-person teams. AI also helped bridge commercial and research-and-development perspectives.

These were bounded innovation tasks. The experiment did not follow the entire journey from an idea to a profitable product, or establish that businesses should dispense with teams. It does suggest that broader access to expertise can improve what an individual brings to a shared problem.

A practical test would ask a product team to examine one opportunity through customer, technical, service and economic perspectives, then have specialists challenge the assumptions. Measure whether the proposal improves and whether fewer unresolved issues reach the next stage. The quality of the subsequent conversation matters as much as the quality of the first draft.

Giving every function an assistant while preserving all its existing incentives may simply make each better at defending its own position. Marketing produces ideas faster, engineering identifies difficulties faster and finance rejects the economics faster. The leadership task is to establish a common problem and settle the trade-offs that faster analysis exposes.

Morgan Stanley makes the workflow tangible

Morgan Stanley's June 2024 Debrief announcement described a connected sequence: with client consent, generate meeting notes and actions, draft a follow-up email for the adviser to edit and send, and save a note into Salesforce. The firm also reported 98% adoption among financial-adviser teams for its separate AI Assistant. That figure did not refer to the newly launched Debrief.

The company account establishes the design and reported adoption, rather than an independently measured financial return. Its practical interest lies in joining the conversation, record and follow-up, with the adviser retaining responsibility. Faster preparation becomes more useful when it reduces the handling around the task.

For another professional-services business, examine what happens immediately before and after the promising application. Re-entering a summary, reconciling versions and chasing missing actions can consume the apparent saving. Follow the improvement through to completed commitments and a more reliable client experience.

DBS connects deployment to reported value

DBS's 2025 annual report describes more than 2,000 models and over 430 use cases delivering approximately SGD 1 billion in economic value during the year. This includes data analytics and AI/machine learning, rather than generative AI alone. The bank also reports that its DBS-GPT assistant supports two-thirds of employees.

These are management-reported results. They show the breadth of one institution's programme without isolating what AI caused or predicting another bank's return from similar tools. Years of investment in information, systems and organisation sit behind the account.

The useful question for a leadership team is how deployment connects to a business result. Which application changed a customer outcome or an economic decision? Which complementary investments made it possible? Give those conditions their share of the explanation, or the next investment case may fund the visible interface and omit what made it work.

The counterexample belongs in the argument

METR's early-2025 randomised study involved 16 experienced developers completing 246 tasks in familiar open-source repositories. Allowing AI increased completion time by 19%. Its February 2026

follow-up found indications that newer tools were helping more, but selection effects and measurement difficulties prevented a reliable estimate of the current benefit. Developers' willingness to work without AI had itself changed.

A disappointing finding deserves attention without becoming a permanent prohibition. A favourable result needs the same discipline. Tools improve, people learn and the tasks they attempt change. Test the actual work, including the effort needed to check and integrate the result, and attach dates and conditions to the conclusion.

Together, these cases justify ambition with discrimination. AI can make expertise more accessible and remove effort around a task. It can also move effort into checking. The decision is which work to change, under which conditions, with what evidence of improvement.

“

Ask what people would have to give up for the change to succeed.

Understand what people are protecting and why

An employee can understand the case for AI and still hesitate to use it. They may be protecting a customer, a professional standard, a position they have worked hard to earn or a little breathing space in an already full week. Calling that resistance is convenient. It tells you very little about what to do next.

Rogers' work on diffusion directs attention to perceived advantage, compatibility, complexity, opportunities to try an innovation and visible results. The COM-B model asks whether people have the capability, opportunity and motivation to perform a behaviour. Both help explain a particular situation. Neither makes a colleague a permanent enthusiast or a permanent obstacle.

Consider three people who have stopped using an assistant. One lacks approved access to the information required. Another found that checking took longer than doing the task directly. The third discovered that finishing early meant receiving more work, while helping colleagues counted for nothing. One adoption score conceals three different management problems.

The first needs an access decision. The second needs a better application, and may be right to stop. The third needs a different agreement about workload and contribution. The diagnosis should leave room for the possibility that the person declining to use AI has understood something the programme has missed.

Sharon's work on leading the human side of change adds a further distinction: knowing what to do and being able to do it under pressure are different achievements. A manager may sincerely support experimentation, then cancel learning time when delivery slips. An expert may endorse sharing knowledge, then reclaim every difficult decision. The contradiction is useful evidence about the pressures and assumptions that need attention.

Ask what the person believes they would lose if the change worked. An expert's value may have rested on being the one who knows. A manager's standing may depend on the size of the team. A conscientious employee may fear submitting anything they did not produce themselves. Those concerns require a conversation about identity and responsibility as well as instruction in a tool.

Understanding them does not remove accountability. It helps make the change possible. An expert can become responsible for the quality of judgement across a wider group. A manager can be recognised for a better service rather than a larger team. Colleagues can learn where personal verification remains essential. The new contribution needs to be real, visible and reflected in how performance is assessed.

This is where commitment begins to differ from compliance. Someone can attend the course and reproduce the approved language while preserving the old method whenever it matters. A more demanding test is whether they can explain the purpose, make a real attempt, discuss a failure and improve the next attempt. Learning has to survive an ordinary working week.

Four conditions reveal whether the bargain holds

Four questions help make these conditions visible. Look at what people believe, what managers permit, what the organisation rewards and who owns the outcome. We propose them as a practical way to diagnose the work, rather than a validated scoring system. Their value lies in exposing a contradiction that someone can resolve.

What people believe

Ask colleagues what they think AI means for their job. Compare their answers with what executives intend. The difference may reflect an unconvincing message, inconsistent local experience or a reasonable concern that has not been answered. The company owes people clarity about its direction, the decisions still open and how they will be involved. In return, it can ask for a willingness to learn and reconsider familiar ways of contributing.

What managers permit

Permission has to exist in the working week. Approved access, protected time and clear boundaries for experimentation make a serious attempt possible. Managers should expect people to try the method and review the result, including failures. If every hour remains committed to existing delivery, the invitation to experiment has little practical meaning. Look at what the manager has stopped as well as what they have encouraged.

What the organisation rewards

Follow the consequences of success. Does faster delivery improve someone's standing, or merely increase their workload? Does teaching three colleagues count in appraisal? Experts need recognition for the judgement they make available to others, alongside responsibility for its quality. A bargain that asks them to share what makes them valuable while rewarding only individual output is unlikely to hold for long.

Who owns the outcome

The top team must decide what happens to the capacity released, and a business owner must act on it. That includes retiring duplicate work and making choices about service, growth or expenditure. Colleagues are more likely to understand the change when they can see what their effort makes possible. Repeatedly asking them to save time without deciding how it will be used weakens the purpose of the whole exercise.

The chief people officer can bring these conditions into one conversation with the chief executive and business owners. Compare roles, sites and teams; combine confidential feedback with

observation of the work. A company average may hide a manager who has solved a problem that others still face. The people function helps make the bargain coherent; operational leaders remain accountable for keeping it.

Follow an employee far enough to see whether the method becomes normal: they understand its purpose, have permission to try, can judge the output, are recognised for contributing and use it without exceptional support. This is a useful way to examine progress, not a fixed sequence everyone must follow. Keep checking whether ordinary use is producing better work and value.

Leaders have to manage their own response

Leaders are participants in the change they are asking others to make. AI may challenge their expertise, status or sense of control too. A senior manager who is unsure of the technology can become impatient with questions, demand unnecessary approvals or retreat behind a specialist. Good intentions do not prevent those reactions from shaping everyone else's experience.

Sharon's work on leading self through challenging change puts attention on the capacity to think clearly and act deliberately when tired, threatened or under pressure. Its practical contribution here is to make self-awareness something leaders practise. Knowing the right behaviour in a workshop is easier than sustaining it when a customer complains and the experiment is yours.

Before a difficult review, notice what you are bringing into the room: fatigue, urgency, defensiveness or a strong need to prove the original decision right. Pause long enough to decide what the situation needs from you. During the discussion, ask for the evidence before defending the programme. Afterwards, seek specific feedback about whether your response made it easier or harder to raise the next problem.

This can be a small, repeatable discipline. A leader might practise listening to a concern without immediately rebutting it, or acknowledging what they do not yet know before making a decision. A trusted peer or coach can help identify the pattern and review progress. The test is visible conduct: a clearer decision, an honest conversation or an error surfaced in time to protect a customer.

Personal resilience also needs sensible working conditions. Asking exhausted managers to become better at coping while adding more demands leaves the cause untouched. Examine their workload, give them space to recover and simplify competing priorities. Leaders remain responsible for their impact, while the top team remains responsible for the system in which they lead.

Where AI supports coaching or personal reflection, agree what information is captured, who can access it and how it may be used. Private developmental conversations should not quietly become inputs to performance assessment. People need a basis for trust before they can speak candidly.

Seven actions that turn AI into performance

Leadership actions must support changed behaviour and value the business can realise.

Leadership action	What leaders do	Test of progress
1 Settle the top team's bargain	Agree outcomes, the employee bargain and what will stop. Align time and incentives.	People understand the change and can see the promises being kept.
2 Redesign the whole journey	Own the customer result, handovers and removal of unnecessary work.	Better customer outcomes, shorter cycles and less rework.
3 Build judgement	Practise on real work. Repeat the checks and behaviours that improve decisions.	Colleagues judge output and handle ordinary exceptions under pressure.
4 Make learning safe and useful	Manage your own response. Review attempts with candour, respect and clear standards.	Errors surface early and the method works without one person rescuing it.
5 Spread through credible peers	Transfer the method and its conditions. Give peers time and recognition for teaching.	A second team sustains the result with ordinary support.
6 Own the economics	Decide where capacity goes. Change the service, budget or deployment of people.	Value appears in customer outcomes or economics without double counting.
7 Keep testing	Compare with a baseline. Check the employee bargain and act on new evidence.	Leaders scale, adapt or stop when the evidence warrants it.

Source: Authors' synthesis of the research and management arguments discussed in this article.

Seven leadership actions turn adoption into performance

These are foundational leadership actions. The specific plan depends on where the organisation is now on the AI value continuum and where it wants to get to. Set priorities and sequence the work accordingly. The actions should reinforce one another, beginning with contradictions senior leaders can remove and interventions proportionate to the change required.

1 Settle the top team's bargain with the business

Agree what AI should improve, which position on the continuum fits and what the company is prepared to change. Name the customer outcome, affected work, business owner and evidence required before expansion. Then explain what that means for employees, including what is known about roles and what remains undecided.

The chief executive must resolve conflicts between the message and the management system. If a services firm rewards only hours charged, faster delivery may penalise the very people who improve it. If specialists gain recognition solely through personal output, teaching others becomes a

charitable activity. Those arrangements may once have served a purpose. Leaders now need to decide whether they still do.

Make learning time real by stopping or postponing something. Remove a report, shorten a recurring meeting or sequence competing demands. Protecting that time during a difficult week is a more credible demonstration of priority than describing the programme as strategic.

Employees should be able to identify what changed in their workload.

The bargain is reciprocal. Leaders provide honest information, workable conditions and support through the transition. Colleagues make serious attempts, share what they learn and accept responsibility for agreed changes. Neither side can promise that every experiment will succeed. Both can be held to the commitments within their control.

2 Redesign the whole customer journey

Choose recurring work with a problem worth solving. A quotation that arrives too late, a support issue that repeatedly returns or a decision delayed by missing information gives the experiment a purpose. Involve the people who do and receive the work before deciding where AI belongs.

Trace the journey from demand to outcome. Find where work waits, information is re-entered and decisions are reopened. Distinguish necessary safeguards from habits inherited from an older system. Some of the improvement may come from removing a pointless step. It remains an improvement even if no new technology was needed.

Take an illustrative industrial-services quotation. Sales gathers requirements, engineering checks feasibility, operations confirms capacity and finance reviews margin. AI could produce the first draft quickly while the customer still waits for those functions to agree. The visible task has accelerated; the queue has not.

A more useful test would combine approved technical information, clear requirements and rules for routing exceptions. Measure time to an accurate, deliverable offer, including rework and the burden on other teams. A fast promise that operations cannot honour is a poor customer experience with a good timestamp.

Give one business owner responsibility for the whole result and a way to resolve disagreements. Technical specialists can build an assistant; the business must decide which commitments it is prepared to make. Including colleagues in that design also lets them identify risks and useful exceptions that a process map may miss.

Retire the old work when the replacement is ready. Temporary duplication may protect a transition, but it needs an owner and an exit condition. Otherwise capable people keep both processes alive, often for sensible local reasons, and the organisation calls the additional effort transformation.

3 Build judgement through repeated practice

A plausible answer is a starting point for judgement. Colleagues need to recognise suitable tasks, missing information, consequential errors and the point at which they should seek help. Capability includes knowing when the tool adds little value. A confident demonstration rarely tests those boundaries.

Practise with real work. Compare a sound output with an equally polished one containing a material error. Ask people to trace the evidence, explain the difference and deal with an ordinary exception. Return to the task after feedback. The next attempt shows more than a declaration of confidence at the end of a workshop.

Sharon's emphasis on small, repeated practices is useful here. Attach one behaviour to a recurring task: check the source before passing on a recommendation, record an override or ask a colleague to challenge an assumption. Review whether it happened and what it changed. A manageable practice embedded in the working day has a clearer route to becoming habitual than a long development plan nobody revisits.

Make review proportionate to consequence. An internal brainstorm and a customer commitment need different checks. Specify who approves, what must be inspected and when a specialist is required. Exhaustive checking of every item may destroy the economics; casual acceptance of consequential claims may destroy trust.

Agree how disagreements between AI and experienced colleagues will be resolved. The accountable person should examine the evidence, explain the decision and record what the outcome teaches. An expert's reasoned override can improve the shared method. An unexplained veto leaves everyone dependent on the expert being present.

An AI agent that can update records or commit resources requires explicit limits on its authority. Establish which actions it may take, who is responsible and how exceptions will be handled. Begin where outcomes can be observed and corrected. The ability to act does not settle the question of permission.

Finally, test the method under delivery pressure. If colleagues routinely omit a check, investigate its clarity, cost and practicality. A standard that works only when the team has spare time is unlikely to be a reliable standard. Change the design or the workload before assuming that another reminder will suffice.

4 Make learning safe enough to be useful

Teams can see consequences that individuals miss. One person's faster draft may become another person's longer checking session. Reviewing the work together makes that transfer visible and gives the team a chance to improve the whole method. The unit of learning should follow the interdependence in the work.

Use a short discussion within the existing operating rhythm. Bring an actual attempt: what improved, what failed, what was hard to judge and what will change next time. Follow up the agreed action. Avoid creating another meeting in which people demonstrate activity to a manager who has no time to examine it.

Edmondson's research connects psychological safety, a shared belief that a team is safe for interpersonal risk taking, with learning behaviour. Here the practical test is how the team responds when someone exposes an error or questions a confident answer. The first awkward disclosure is a leadership moment. Your response helps determine whether the next problem is brought to you.

Combine respect with clear standards. An honest mistake within an agreed experiment calls for correction and learning. Concealing a known defect requires a different response. Make that distinction explicit, and hold the same standard when the mistake belongs to a senior colleague. People notice exceptions made for rank.

Check whether the capability is shared. If one enthusiast rescues every attempt, ask another colleague to complete the task, explain the checks and handle an ordinary exception. The difficulty they encounter identifies what needs teaching or simplification. It is more useful than another polished success story.



The first awkward disclosure is a leadership moment.

5 Spread the method through credible peers

Willing teams with worthwhile work are a sensible place to begin. Once the method shows promise, try it where conditions are less convenient: a busy manager, uneven information or colleagues with less confidence. That is where you discover what the method needs in order to travel.

The second team provides a particularly useful test. Can it reproduce the result with ordinary support? If success depends on the original enthusiast or unusually generous attention, that dependence belongs in the investment case. Publicising the pilot more widely does not remove it.

Choose peer practitioners for patience, judgement and the trust of colleagues. Give them time to help and recognise the contribution in appraisal. Their task includes listening to objections and demonstrating recovery from mistakes. People need to see someone credible working through the difficulty, as well as displaying the finished result.

Transfer the whole method: its purpose, inputs, checks, common failures and demands on adjacent teams. Include the conditions that allowed it to work. Where the receiving manager cannot protect practice time or reward teaching, address that gap before concluding that the team lacks enthusiasm.

Keep essential standards common while allowing useful local adaptation. Customer needs and systems differ. Insisting on identical steps everywhere may undermine the service; letting every team reinvent everything may undermine the economics. Be clear about which decisions belong locally and how improvements return to the shared method.

Ask whether the receiving team still uses it after intensive support ends and can help a third team learn. When the method becomes an ordinary part of delivering good work, sustained use starts to mean something. Attendance at a launch establishes much less.

6 Give released capacity an economic owner

An hour saved needs somewhere useful to go. Fragments of time may reduce overload or improve quality without reducing expenditure. Those outcomes can matter greatly. Describe them accurately, then make a separate decision about whether capacity can be sold, redeployed or removed from the cost base.

Consider an illustrative calculation. One hundred colleagues each release two hours a week over 46 weeks: 9,200 hours of potential annual capacity. At an assumed loaded cost of £50 an hour, that has a salary-cost equivalent of £460,000. The payroll has not changed. Calling the figure a cash saving would conceal the next management decision.

Suppose a quarter of those hours can replace external support costing £80 an hour. Avoided spending would be £184,000: 2,300 hours multiplied by £80. Subtract annual tools and support of £60,000 and one-off implementation of £40,000, and the first-year net benefit is £84,000, provided the external spending disappears and quality holds. All figures are illustrative assumptions, not a client result or forecast.

Finance and the business owner should examine when and where the time is released, whether contracts or staffing can change and whether another team has absorbed extra work. If the plan is growth, establish the demand. Do not count the same capacity once as lower cost and again as additional revenue.

Also consider what becomes commercially possible. A service previously too expensive to provide may become viable. Existing customers may receive support the company could not afford to offer. Test the need, willingness to pay and contribution after the full cost of delivery. Better economics can change the proposition as well as the cost of producing it.

“

The payroll has not changed. The next management decision still matters.

The business owner must turn that choice into action. Finance tests the economics; technology supports the method; the people function helps redesign roles and manage the transition. All are necessary. Responsibility for changing the service, budget or deployment of people must still have a name attached to it.

7 Keep testing and act on what you learn

Treat a useful AI method as a proposition with evidence behind it. Record the task, people, tool version, checks and result. Another team can then judge whether it fits, and the original team can recognise when a change in conditions makes yesterday's conclusion unreliable.

Establish the baseline before claiming improvement. Follow quality, rework, elapsed time and effort across the workflow. Random allocation or a well-designed staggered introduction can strengthen the comparison where feasible. Elsewhere, combine operational measures with checked samples and acknowledge what cannot be attributed confidently to AI.

Listen to employee experience without treating it as the productivity result. Feeling quicker and finishing sooner are different things. So are satisfaction with a tool and better service for the next team. Use feedback to locate friction, then follow the work to establish its consequences.

Look for the conditions that sustain the result. Is practice time still protected? Are colleagues reporting errors? Has the expert become an exhausted source of unofficial support? A short-term gain accompanied by rising overload deserves attention before the organisation declares the application ready to scale.

Be willing to stop. An application that does not justify its full cost is consuming attention that could go elsewhere. Record why it failed, so a material improvement in the technology or the work can justify another test. METR's evolving findings illustrate why neither optimism nor disappointment should become a permanent operating policy.

Gabriella Mirabelli, of Amaru Black, proposes that leaders may lower their expectations after disappointing implementation while leaving governance unchanged. This remains an unvalidated hypothesis. It nevertheless raises a useful question about whether the experiment tested the technology's capability or the limits of the organisation's willingness to change.

Distinguish a weak answer from a sound recommendation that nobody has authority to implement. The first may require a better model, information or task. The second requires a management decision. Both can appear in the same disappointing benefit report.

Give the review authority to scale, adapt or stop. Agility means shortening the time between discovering something and changing what the business does. If the review can only commission another report, the organisation has retained the delay while improving its description.

Make the first 90 days produce decisions

Use roughly 90 days as a first management cycle, with timing suited to the work. It is long enough to demand progress and short enough to expose drift. It is not a promise that every application will establish its economics within a quarter.

In the first month, the chief executive and top team should settle the intended outcomes and employee bargain. Ask the chief people officer to examine what people believe, managers permit and the organisation rewards. Select two or three recurring workflows, name their owners and establish the baseline. Stop enough existing activity to make room for a serious test.

During the second month, practise the revised method with the people who do the work. Use experts to define quality, give managers support with difficult conversations and review the effect on customers and adjacent teams. Check whether the promises about time, access and recognition are being kept. A missing permission or conflicting target is now a leadership action, rather than a survey finding.

Use the final part of the cycle to test transfer and economics. Ask another team to reproduce the method with ordinary support. Establish what happened to released capacity and which old work can end. Expand where the evidence warrants it; specify the next learning question where it does not. Stop applications whose value fails to justify their cost.

The review should leave the chief executive able to explain what improved for the customer, which management obstruction was removed and where the capacity went. The chief people officer should be able to show whether the employee bargain was understood and honoured. Together, they should identify the evidence that changed a decision, including one of their own.

The leadership choice comes back to us

Cheaper intelligence creates room to do work that was previously impractical, expensive or slow. The organisation must decide what to make of that opportunity. It can accumulate more output and more checking, or redesign the work so that customers and colleagues experience a worthwhile difference.

That change asks something of people. Experts may have to find new ways to contribute. Managers may have to give up familiar controls. Employees may have to learn while facing uncertainty about their future. Leaders owe them honest information, respect and conditions in which a serious attempt is possible. That is sound management, and it is also how people deserve to be treated.

Our judgement is that the ability to learn and reorganise around widely available technology will become an important source of advantage. It requires leaders who can expose inconvenient results, keep difficult promises and change arrangements they once defended. Some of the work will involve substantial redesign. Some will involve stopping a meeting and allowing an expert to teach.



Choose one worthwhile change, and make it possible for people to deliver it.

At the next investment review, take one workflow and follow the argument all the way through. What must change for the customer? What must people learn or relinquish? Which decision are you still asking them to work around? Who will turn the improvement into a business result?

Choose one worthwhile change, and make it possible for people to deliver it. When the new method survives ordinary pressure, improves the customer's experience and produces value someone can account for, intelligence has begun to pay.

Sources and further reading

Research findings apply to their studied settings. Company figures are reported outcomes. Discussion papers and conceptual frameworks inform the argument without establishing forecasts of AI returns.

Price, C., and Toye, S. (2017). *Accelerating Performance: How Organizations Can Mobilize, Execute, and Transform with Agility*. Wiley. [Source](#)

Keller, S., and Price, C. (2011). *Beyond Performance: How Great Organizations Build Ultimate Competitive Advantage*. Wiley. [Source](#)

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative AI At Work. *The Quarterly Journal of Economics*, 140(2), 889–942. [Source](#)

Dell’Acqua, F., et al. (2026). The Cybernetic Teammate: A Field Experiment On Generative AI And Teamwork. *Organization Science*. [Source](#)

Morgan Stanley (2024). Morgan Stanley Wealth Management Announces Latest Game-Changing Addition To Suite Of GenAI Tools. 26 June. [Source](#)

DBS Group Holdings (2026). Annual Report 2025: Letter From Chairman & CEO. [Source](#)

Becker, J., Rush, N., Barnes, B., and Rein, D. (2025). Measuring The Impact Of Early-2025 AI On Experienced Open-Source Developer Productivity. *METR*, 10 July. [Source](#)

Becker, J., et al. (2026). We Are Changing Our Developer Productivity Experiment Design. *METR*, 24 February. [Source](#)

Rogers, E. M. (2003). *Diffusion Of Innovations*. Fifth edition. Free Press. [Source](#)

Michie, S., van Stralen, M. M., and West, R. (2011). The Behaviour Change Wheel: A New Method For Characterising And Designing Behaviour Change Interventions. *Implementation Science*, 6, 42. [Source](#)

Lawson, Dame Emily, and Price, C. (2003). The Psychology Of Change Management. *McKinsey Quarterly*, 1 June. [Source](#)

Edmondson, A. (1999). Psychological Safety And Learning Behavior In Work Teams. *Administrative Science Quarterly*, 44(2), 350– 383. [Source](#)

Stanford HAI (2025). AI Index 2025: State Of AI In 10 Charts. 7 April. [Source](#)

Klarna Group plc (2026). Annual Report On Form 20-F For 2025. [Source](#)

Mirabelli, G. (2026). AI Adoption Framework: Core Hypothesis And Theoretical Foundation. Version 4.3, 30 September. Amaru Black LLC. Unpublished discussion paper.

Toye, S. (2026). Leading Self Through Challenging Change. Version 2.0, 20 August. Unpublished development paper.

Toye, S. (n.d.). Leading The Human Side Of Change. Leadership Mindset Partners. Unpublished discussion paper.

AI-Culture Framework Prototype (2026). AI, Culture And The Chief People Officer. September. Unpublished conceptual prototype.

About the authors



Colin Price

Organisation performance

Colin Price advises boards, chief executives and investors on organisational performance through Accelerating Performance Consulting. He is Chair of itim Group plc and was previously Chair of Co-Company. A former McKinsey Senior Partner, he led its worldwide Organisation Practice for seven years. His executive roles have also included CEO of Heidrick Consulting and COO of Quintet Private Bank.

His work connects strategy, organisational health and leadership behaviour to business results. He is the author or co-author of seven books, including *Accelerating Performance*, with Sharon Toye, and *Beyond Performance*, with Scott Keller.

colin@acceleratingperformanceconsulting.com



Sharon Toye

Leader and team performance

Sharon Toye is a founding partner of Accelerating Performance Consulting and advises chief executives and senior teams on leadership, team effectiveness and lasting behaviour change. She combines commercial experience with psychological insight as a Master Certified Coach with the International Coaching Federation and a UKCP-registered psychotherapist.

Previously, she was a founding partner of Co-Company and led the global Team Acceleration practice at Heidrick & Struggles. She holds an MBA from London Business School and master's degrees in organisation psychology and integrative psychotherapy. She is co-author, with Colin Price, of *Accelerating Performance*.

sharon@acceleratingperformanceconsulting.com

Intelligence Pays is part of our series on leadership and organisational performance. All papers are free to read at acceleratingperformanceconsulting.com/read

© Accelerating Performance Consulting Limited. Registered in England and Wales, company number 12224907.