


Where Judgement Will Live?

What AI implementation
does to the capability
it depends on.

David J. Norris

The Brain Commons • Notes on Human Agency • Essay 03



Where Judgement Will Live?

What AI implementation does to the capability it depends on

David J. Norris

Author of *The Brain Commons*

David J. Norris is an Australian author and senior occupational therapist, and the author of *The Brain Commons* — an inquiry into the shared conditions that allow people, communities and institutions to remain capable of attention, judgement, relationship, participation and care. Across more than twenty-five years in neurological and cognitive practice, he follows that thread past the clinical frame in his publication, *Notes on Human Agency*.

A paper for boards, executives and the functions that approve AI-enabled redesign of work: strategy, workforce planning, technology, risk and continuity.

PUBLICATION NOTICE

Where Judgement Will Live?

What AI implementation does to the capability it depends on

David J. Norris
The Brain Commons

Publication
August 2026

© 2026 David J. Norris.
All rights reserved.

This publication may be shared in its complete and unaltered form for non-commercial purposes, provided that the author and publication are clearly attributed. No part may be reproduced, adapted, republished or used commercially without prior written permission, except for brief quotations used with appropriate attribution.

The analysis and views expressed in this paper are those of the author. The publication is intended to contribute to discussion about AI adoption, human capability, work and organisational governance. It is provided for informational and professional discussion purposes and does not constitute legal, financial, regulatory, clinical or other professional advice.

References to external organisations, standards, frameworks and published sources are included for analysis and discussion. Their inclusion does not imply endorsement by those organisations, nor endorsement of this publication by them.

Suggested citation

Norris, D. J. (2026). *Where Judgement Will Live? What AI implementation does to the capability it depends on*. The Brain Commons.

Permissions and enquiries

David J. Norris
The Brain Commons
DavidJNorris.com

The numbers can all look good. They say little about what has happened to the work.

Processing time falls. Output rises. Staff spend less of the day on work nobody wanted to do in the first place. A business case that began as a projection starts returning figures that hold up under scrutiny, and an implementation that was an experiment eighteen months ago becomes simply how the work is done. There is every reason to call this a success, and in most respects it is one. But a dashboard reports what was measured, and what was measured was the thing the organisation set out to change. It is quieter about who now sees the difficult cases, about where review has replaced production, about which judgements have moved to another desk or another pay grade, and about what the people coming up behind no longer encounter often enough to learn from.

These consequences are shaped by decisions made in different rooms. Strategy identifies the opportunity. A business case prices the gain. Workforce planning models roles and headcount. Technology teams configure and integrate. Change programmes establish new ways of working. Risk looks for the ways it might fail. Every one of those functions can reach a defensible decision within its own terms, on its own evidence, against its own criteria — and the cumulative effect on judgement will still be nobody's paper. Organisations do not fail to plan AI adoption. They plan it exhaustively. The redesign of work is divided among functions that are never required to see it whole, and no single approval in the sequence is an approval of the thing that actually happened.

AI-enabled redesign can remove work outright, but it can also redistribute it. People may do less of the ordinary run of work and more of what begins when the ordinary run breaks: checking, interpreting and resolving the cases that do not fit. A sequence of bounded tasks can become oversight across several streams at once. The saving is booked where the time was taken out. The demand that replaces it can surface somewhere else entirely — another role, another team, another cost centre, another year's operating conditions — where nobody is looking for it, because nothing there was supposed to have changed.

Consider an illustrative claims operation in which the exception rate falls from ten per cent to four. On any reasonable reading that is a good result, and it is the result the system was bought to deliver. Now suppose volume quadrupled over the same period. The rate has improved by well over half. The people handling exceptions are dealing with sixty per cent more difficult cases than

before, and difficult cases are not merely more numerous — they are what the working day now consists of, one after another, with the ordinary work that used to sit between them gone. Both numbers are true. Only one of them is going to the board.

Time released does not, by itself, tell us how much usable capacity has been created. In June 2026, NHS England announced that 505,000 staff would be given Microsoft 365 Copilot, following a trial across around ninety organisations which reported an average saving of forty-three minutes per person per day. The announcement translated that figure into roughly two days a month and five weeks a year, and connected the rollout to more time with patients, greater productivity and shorter waits. Those are different kinds of claim. The time conversion is arithmetic; the service consequences depend on what the released time becomes inside the work.

Because minutes do not interpret themselves. They might become time with patients. They might become more patients. They might become room for the review and supervision that had quietly been squeezed out of the week, or the half hour in which somebody finally thinks properly about a case that has been bothering them. They might equally become the new baseline expectation, absorbed into the roster before anyone has decided anything at all. The technology does not settle this. Rostering settles it, and targets, and workflow, and the judgement of a local manager under pressures of their own. Two organisations can buy the same system, record the same saving, and produce entirely different working lives.

Alongside assisted performance there is a second number, rarely asked for: what remains possible when the system is unavailable, degraded, withdrawn, or confidently wrong. The gap between the two is not in itself evidence of a problem — it may be precisely the gain the organisation paid for. But left unmeasured, it makes dependence, retained capability and recovery impossible to tell apart. Nobody works unaided, and the point is not self-sufficiency. It is knowing which of the supports an operation runs on are held inside it, and which are held by somebody else, on terms set elsewhere and revisable without notice. The concern has a long lineage in automation research, and recent safety-critical guidance makes skill erosion, retained operational competence and fallback operation explicit. Across the reviewed corpus, I did not find a routine requirement to measure what people can still do when the AI system is unavailable or degraded, or to connect that capability to a particular implementation across time.

No organisation begins with a blank workforce. People arrive carrying judgement formed somewhere else — by previous employers, by a profession, by a supervisor who sent the work back, by years of ordinary practice alongside people who knew more than they did. Much of that formation was not separately funded as formation, because it happened inside the work: in the second draft, the checked figure, the case discussed after the meeting had broken up, the awkward problem handed deliberately to somebody who was not quite ready for it. Change the

task and that formative function can disappear with it, silently, while output holds or improves. Capturing what experienced people know does not resolve the problem. It defers it, and leaves a harder question behind: what will form the next person whose knowledge is worth capturing?

Many contemporary AI governance arrangements rely on human oversight. The competence that oversight requires is increasingly specified, but the reviewed corpus does not require an implementation to account for how that competence will be retained or formed over time. The wider version has the structure of a commons: capability formed across organisations and professions, while the organisations drawing on it are not always the ones that bear the cost of replenishing it.¹ The version that concerns this paper is nearer to hand, and more tractable. Not the fate of a profession across twenty years, but a decision taken in a particular organisation, on a particular Tuesday, by people who could have asked a different set of questions and were never required to.

The paper therefore stays with a small number of questions that should be answerable before a substantial AI-enabled redesign is approved, but that currently sit across several organisational functions. What human work is changing, and what else was that work doing? Where will the displaced demand travel, and who will be holding it when it arrives? What will occupy the capacity released — and has anyone actually decided, or will the roster decide by default? What must remain possible when the system cannot carry the load? And who, once the implementation team has been redeployed and the saving has been booked, is still asking whether the judgement this organisation depends upon is being formed, exercised and renewed?

Where Judgement Will Live? sets out how those questions can be put in practice — inside a business case, a procurement, a workforce plan, a paper going to a board — and why they currently fall between arrangements built for other purposes. The gains a business case is designed to capture can appear quickly. Consequences for judgement may emerge later, after the people who approved the original decision have moved on. The people doing the work will still be there.

Where Judgement Will Live?

A man I saw years ago had had a stroke and did not believe it. Ten minutes in his company and you would have thought nothing had happened. I asked him to draw himself before and after. He could not do it, and could not say why. His wife had watched him brush his teeth for eight minutes, the same movement up and down, until she asked whether he might be finished.

He had always cooked. He wanted to make a beef stroganoff, so I asked him first how he would go about it and how long it would take. Forty-five minutes, he said, and he listed the steps correctly. Then I let him go.

He opened the fridge and left it open. He put oil in the pan and the heat on. At forty minutes he was still chopping the steak, which had long since become mince, and the pan behind him had begun to smoke. I turned off the stove. The fridge was still beeping.

I asked whether there was anything in the room that needed his attention. Not the fridge. Anything. He looked around and could not find it. When I named it, he crossed the kitchen and shut the door without difficulty.

Everything in that room was giving him feedback. The beeping. The smoke. The meat going to paste under his hands. None of it reached him. The friction was all there. What had been damaged was the part of him that could receive it.

What eventually worked was a score. A screen that told him how he had done and asked him to do better. He came back to it because he wanted to beat it. It was the first time I saw a technology close a feedback loop that a person's own attention no longer could.

He could not see the problem until somebody outside the room named it. That is the part I have thought about since.

Twenty-five years later I see a related omission in the way organisations redesign work. It appears across business cases, workforce models, technology implementations and change programmes. Each asks a legitimate part of the question. What can be automated. What time can be removed. Which roles must change. What people will need to learn. What adoption will require.

What is rarely held alongside those decisions is what cognitive work is moving, what capability the old work was sustaining, and where the resulting demand will now be carried. The task appears in the model. Its other functions may not.

Every substantial AI implementation that changes how work is done also changes some part of an institution's cognitive architecture: where judgement sits, who holds which knowledge, how decisions travel, and who is positioned to notice when something has gone wrong. That change is not made by the technology. It is made in the choices about what to automate, what remains human, what gets escalated, what must be verified, how fast work arrives, and what happens when the system is wrong or unavailable. Two organisations can install the same system and arrive at very different consequences.

Adoption is currently governed as a matter of technology, productivity, risk, privacy and cost. The same redesign can also redistribute judgement, move demand into new places and alter how capability is formed. Those consequences are not routinely held together as a single governance object.

What judgement means here

Before going further, I need to be more precise about the word at the centre of this paper.

Judgement is used loosely in discussions of artificial intelligence. It can mean prediction, classification, recommendation, choice, professional discretion or simply the answer that appears at the end. Those are not the same thing.

As I use it here, judgement is the situated work of recognising what kind of situation one is facing, determining what matters within it, weighing evidence, consequences and obligations, and deciding what response is warranted. That response may be to act. It may also be to wait, verify, seek another view, refuse, or escalate. A decision is the eventual commitment. Judgement is the work through which that commitment becomes warranted.

Judgement therefore draws on cognitive capability without being reducible to it. Attention, memory, inhibition, flexibility and experience matter because they affect what a person can perceive and hold in mind. Knowledge matters because it gives the situation a history and makes patterns recognisable. Metacognition matters because it allows somebody to recognise uncertainty and the limits of what they know. But none of those capacities, separately or together, settles what should be done. Judgement begins where facts, uncertainty, values and consequences have to be brought into relation in a particular case.

Nor does all of that work occur inside one person. Judgement is assembled through colleagues, records, instruments, procedures, professional traditions, prior decisions and the conditions of the situation itself. A nurse notices something because another nurse taught her what to look for. An engineer questions a figure because it does not fit the behaviour of a system she has encountered before. A team reaches a better conclusion because somebody has standing to disagree. Technology has always participated in these arrangements. AI becomes another participant, and potentially a consequential one.

This complicates the title of the paper. Judgement does not move as a single object from a person into a machine. Its constituent work can be distributed. A system may detect the signal. A model may classify the case. A worker may interpret the recommendation. A manager may hold the authority to act. The organisation may carry the legal consequence. It is possible for judgement to be described as human because a person remains formally accountable, even though that person no longer has the information, time or practical authority needed to challenge what the system has produced.

The question is therefore functional before it is metaphysical. This paper does not need to settle whether a machine can truly judge. It asks how the work through which judgement becomes possible has been arranged: who or what can notice, interpret, contest and act; where authority sits; how competence is formed; and who remains answerable for what follows.

Reliable judgement does not mean infallible judgement. It means judgement sufficiently accurate, calibrated, context-sensitive, contestable and robust for the consequences it carries. A conclusion may be accurate in an ordinary case and brittle in a novel one. It may be technically correct but reached through an arrangement nobody can question. It may be consistent while repeatedly excluding something the system was not designed to see. Reliability belongs not only to the answer, but to the conditions through which the answer was produced, challenged and revised.

None of this assumes that existing human judgement is always superior. Human practice can be inconsistent, biased, tired, hierarchical and badly calibrated. Technology can remove poor work, widen access, improve detection, make neglected evidence visible and support decisions that would otherwise be worse. The proper comparison is not between an idealised human and an isolated machine. It is between the working arrangement before implementation and the human–AI arrangement that follows it.

Nor should every existing skill be preserved. The disappearance of a task is not inherently a loss, and difficulty does not deserve protection for its own sake. What must remain available will depend on the work: the seriousness and reversibility of its consequences, the likelihood of novelty or system failure, the expertise required to detect an error, the availability of alternative support, and

the time it would take to recover capability once lost. The aim is not maximum human self-sufficiency. It is sufficient, accessible and renewable capability for the obligations the organisation continues to carry.

This is why cognitive offloading cannot be classified as either gain or loss in advance. Removing mechanical work may create room for attention, reflection and more consequential judgement. It may instead increase throughput, concentrate exceptions, shift people from active production into passive verification, or remove the encounters through which later expertise would have formed. Released effort has no inherent destination. The technology does not decide whether it becomes room to think or simply a new expected rate of work.

Repetition is not inherently formative either. Ordinary work develops judgement when it contains opportunities to notice variation, attempt the task, receive feedback, correct an error, observe somebody more experienced and carry increasing responsibility. Repetition without those conditions may do little more than reproduce existing habits. The concern is therefore not that routine work is disappearing. It is that learning, spacing, feedback and exposure can disappear with it without being recognised as part of what the task was providing.

AI implementation can therefore redistribute more than effort. It can redistribute perception, knowledge, authority, learning and responsibility. Assisted performance may improve while some of those conditions strengthen and others weaken. The governance task is to make that redistribution visible before the new arrangement becomes ordinary.

That requires looking backwards before looking forwards. Before asking what new capacity AI might create, we need to understand what the work selected for removal was already carrying. A task may produce an output, but it may also space difficult decisions, expose somebody to normal variation, make an exception perceptible, create an occasion for correction, or allow a less experienced person to work beside somebody who knows more. Those functions rarely appear in a process map, role description or business case. They are nevertheless part of what is being redesigned.

What the work was doing

Start with what happens when a layer of an organisation goes. The removal may be attributed to AI, accelerated by it, or merely justified by it. Organisations have been flattening for decades without any of it. The cause matters less than the accounting. When a layer goes, it has usually been costed accurately. What has not been costed is what that layer was doing when it was not doing its job description.

A senior nurse reads a new policy and works out what it means on a ward that is short two staff, with a bay of patients the policy did not anticipate. She writes nothing down. She adjusts how the shift runs, tells three people what to watch for, and the policy becomes practice without anyone recording the translation. Her job description says clinical care and supervision. It does not say she is the mechanism by which institutional intention survives contact with a real ward. It does not say that the junior nurses on that shift are learning, by watching her, how to do the same thing in ten years.

Translating between the people who decide and the people who deliver. Noticing that a figure is wrong before it reaches a board. Absorbing a complication so completely that it never becomes an incident. Explaining, on a Tuesday afternoon, why we do it this way here. None of that appears in the role description. So none of it appears in the saving.

Automation can remove a task while leaving work around it. That work changes speed, location and form. A first draft may disappear while the judgement needed to review it remains. Routine reconciliation may shrink, leaving people with the exceptions, discrepancies and cases that refuse to settle cleanly. A sequence of bounded tasks becomes oversight across many streams at once. The saving is counted at the point where the time was removed. The new demand disperses across the system, and often has no single cost centre. Automation can remove much of the routine run first. What remains for people may be the edge: the ambiguous case, the exception, the thing that needs another look.

The room around them

Imagine a claims operation. Ten thousand claims a year, one in ten needing a decision that is genuinely difficult. A thousand hard cases, spread across twelve months and across a team, sitting alongside nine thousand that were largely mechanical.

That last part is easy to pass over, and it is the part that matters. The hard cases arrived with room around them. A run of straightforward files after a bad one. An hour of routine in which the difficult decision made that morning could settle. Work a newer colleague could be handed without much supervision, and which you could look up from when they asked. The nine thousand were not only output. They were the spacing.

The system arrives and takes the mechanical work. Volume rises, because now it can — the marginal cost of processing a claim has fallen, and the organisation does what organisations do with that. Forty thousand claims. And the exception rate falls, genuinely, from one in ten to one in twenty-five.

BEFORE

10,000 claims · 1 in 10 difficult

1,000 hard cases a year, arriving with room around them.

AFTER

40,000 claims · 1 in 25 difficult

1,600 hard cases a year. The rate improved by more than half; the human load rose sixty per cent.

Four per cent of forty thousand is sixteen hundred.

The rate improved by more than half. The number of difficult decisions arriving at a person rose by sixty per cent. And the nine thousand ordinary files that used to sit between them have gone, so the sixteen hundred now come one after another, in a day with no shallow end.

Nobody lied. The exception rate really did fall, and the people who wrote the business case were describing something true. The people now working sixty per cent more hard cases in the same hours are also describing something true. There is no document in that organisation where those two facts meet.

The caveats are considerable, and I would rather put them here than bury them. The numbers above are invented — one arithmetic, in one function, chosen to make a mechanism visible rather than to report a finding. Everything turns on volume rising. Hold volume flat and the same system delivers four hundred hard cases where there were a thousand, which is straightforwardly good news, and the business case is simply right. So the claim is conditional, and the condition is empirical: does throughput rise when unit cost falls? Often. Not always, and not by any fixed amount.

Then a real decision, announced while I was writing this. In June 2026 NHS England said it would give 505,000 clinicians and support staff access to a general-purpose AI assistant, following a trial across more than thirty thousand staff in around ninety organisations. The trial found the tool could save an average of forty-three minutes per person per day, which the announcement rendered as roughly two days a month, or five weeks a year.²

Take the figure at face value. The announcement names several intended destinations for the released time: more time with patients, greater productivity and faster treatment. What it cannot settle is how forty-three minutes will actually be used inside a local service. That will be decided by rosters, targets, vacancy rates, workflow and the ordinary pressure of demand. The technology releases time; the work system decides what fills it.

TABLE 1

Time saved, and what it says about capacity

An improving rate and a rising human load are not in conflict, and a figure describing time is not a figure describing capacity.

10% → 4% Exception rate, before and after. A genuine improvement of more than half.	+60% Difficult cases now landing on a human, in the same hours. Illustrative arithmetic.	43 min Potential reported daily saving per person in the NHS England trial. What arrives into it is not settled by the technology.
---	--	--

Sources: exception-rate and hard-case figures are the illustrative claims-operation arithmetic set out above, not observed data. The 43-minute figure is the reported average daily saving per person in the NHS England trial of more than 30,000 staff across 90 organisations (see note 2).

Before asking what the released time achieved, ask what the removed work was doing. The first question looks forward. The second looks back, and can be harder to answer once the work itself has already changed.

Softer ground

Beyond the arithmetic I am on softer ground, and I want to be clear about where the ground changes.

That automation removes one class of work and leaves people with a newer one — monitoring, exception handling, recovery — is among the oldest findings in the field.³ That demand migrates in that direction here follows from what these systems do. That much I would defend.

Past that, I am describing a pattern practitioners recognise rather than one that has been measured. Fewer people left in the middle to absorb complexity. Senior accountability intensifying because there are fewer shoulders for it. Junior formation thinning as the tasks that did the forming disappear. An operation becoming faster and more brittle at the same time.

Test them against your own organisation. If they are wrong where you work, that is worth knowing, and I would rather be corrected than agreed with. I did not find an instrument in the reviewed corpus that routinely collects the answer either way.

WHAT IS CLAIMED, AND WHAT IS NOT

I would rather set this out plainly than leave a reader to work out for themselves how much weight each part of the argument will bear.

ARGUED, AND SUPPORTED BY THE MATERIAL REVIEWED

- AI implementation can redistribute cognitive demand rather than simply remove it.
- Local measures can improve while the absolute burden of difficult human work rises.
- Assisted performance and retained capability are different things, and need governing as different things.
- In the instruments reviewed in full, I did not find a requirement for workforce capability measures to be read against a specific AI implementation across time.

OFFERED FOR TESTING, NOT AS FINDINGS

- That middle layers are thinning in ways that matter.
- That difficult judgement is concentrating among fewer people.
- That junior formation is being altered as routine work disappears.
- That organisations may become faster while losing some of their capacity to absorb complexity and recover from it.
- That rapid work redesign can outpace the formation, distribution and renewal of the judgement the redesigned system continues to require — referred to here provisionally as *judgement shock*.

These are recognisable in practice. They are not established by the evidence reviewed here, and the instruments examined do not routinely collect what would be needed to test them.

NOT CLAIMED

- That AI causes these effects.
- That difficulty is worth preserving for its own sake.
- That organisations never examine independent capability.

Technology can remove poor work, widen access, support learning, and create capability that did not exist before. What follows depends substantially on how the work around it is designed.

Strain that does not announce itself

People may be tired. That is worth attending to, but it is not the interesting part. The harder case is the person who knows exactly what to do and is too stretched to do it well, which is not a failure of competence and should not be handled as one. It can happen when more of an operation comes to rest on fewer people. Queues may lengthen. Escalations may rise. Something that once took a day can begin taking a week, without anyone being able to say quite when the change started. All of it can appear well before anyone in the building would describe themselves as overwhelmed, and without anyone connecting it to a system that went live eighteen months earlier. Whether the part of the operation now carrying that work can absorb it is an operational question with an operational answer. The instruments reviewed do not ask it.

Redistribution is not the argument. It is the mechanism by which the argument's subject changes. People carry on working while the conditions around them shift. The work that taught them. The colleague they could ask. The half hour in which they might have noticed something odd about a case. Any of these can go without it ever surfacing as a decision about judgement.

How the loss can hide itself

The man in the kitchen is an analogy, not a model. An organisation is not a person with an injured faculty, and I do not want to push the comparison further than it will go. But the two share one feature, and it is the feature that ought to worry us: the loss can damage the very arrangements through which it would otherwise become visible.

At organisational scale the reasons for that are dull rather than sinister. Measurement is built around throughput, because throughput is what the business case promised, and nobody funds a measure of something they were not expecting to change. Incentives reward the saving in the year it is booked. The consequences may arrive later. People move on. And information about strain can travel upward as sentiment rather than as a number — a comment in a staff survey, a retention figure nobody can quite account for, a manager saying the team is a bit under it at the moment — which makes it easy to hear and easier to discount. With no recognised object to attach it to, the question falls between agendas rather than being refused on any of them.

Effort that looks like waste

At the individual scale something quieter happens. Some of the effort by which capability forms is easy to mistake for waste while it is going on.

Not all struggle is formative, and I want to be careful here, because the opposite claim is sentimental and does real harm. A great deal of difficulty in modern work is poor design, absent support, or repetition that teaches nobody anything, and removing it is simply good. But the effortful part of learning is often indistinguishable, from the inside, from inefficiency. It feels like being slow. So when a system offers to take it away, the offer is welcome.

The man in the kitchen was given a screen, and the screen closed a developmental loop that his own attention could not. Technology can also support formation, deliberately and well. Which leaves a narrower question, and a better one. When work changes: what developmental functions are going, what is replacing them, and will the replacement form the capability we are going to be depending on later?

The second number



Two numbers. What a person can do with the system, and what they can do without it.

The second tells you what remains when the system is unavailable, degraded, withdrawn or simply wrong. Neither number matters more than the other. The gap between them is not in itself a problem — it may mean the system is doing exactly what it was bought to do, or that somebody is well supported in work they could not otherwise attempt, which is usually the point of buying it. It may also mean dependence. Or a capability that was never formed in the first place. Or a risk that shows itself only on the day the system stops.

The gap will not tell you which. What it does is make the question available. Hold only the first number, and there is nothing to ask it of.

Whose support is it

Nobody works unaided. The nurse has three colleagues watching for what she asked them to watch for. The surgeon has a checklist. The accountant has a standard, and every profession runs on notation somebody else invented. This paper has already argued that capability is never privately formed. So why draw the line at this system and not at the others?

The better question is not where a support sits but how much say the people and the organisation retain over it. Colleagues leave. Standards change. The line is not simply inside and outside. But an experienced colleague, a working protocol, a habit of practice — these are supports the organisation is party to, can watch changing, and can go some way towards replacing. A licensed system is supplied on terms set somewhere else, and can change without anyone in the building having decided anything at all. Nobody is being asked to work alone. The question is more ordinary than that. When a support changes, fails, or goes away without warning, what can the people doing the work still carry, and what can the organisation actually replace?

Where the problem is already recognised

The concern already appears explicitly in safety-critical guidance. Joint guidance on artificial intelligence in operational technology, published in December 2025 by the American and Australian cyber-security agencies with partners in five other countries, names skill erosion as an operational risk in its own right.⁴ Heavy reliance on the system, it says, may cost personnel the manual skills they need when it fails. It recommends that teams maintain operational competencies alongside the AI, that failsafe mechanisms revert AI-enabled processes to traditional automation or manual operation, and that thresholds be set for defaulting back. Voluntary guidance rather than requirement, and written for dams and substations rather than for offices — but the logic is there, and it is explicit. Business continuity standards separately treat exercising and testing as how an organisation finds out whether its arrangements work when they are needed.⁵

Across the public instruments reviewed in full, I did not find a routine requirement to measure what people can still do when the AI system is unavailable or degraded. Enterprise and national adoption surveys track use and reported productivity. Human-capital reporting has moved closer to this territory. ISO 30414:2025 now contains requirements as well as recommendations and retains reporting across productivity, mobility and succession planning, and skills, capabilities and development. ISO/TC 260's public account of the revision also identifies artificial intelligence and workforce upskilling. I did not obtain the full normative text, so no absence claim in this paper rests on ISO 30414:2025.

The narrower finding is this: in the corpus reviewed in full, I did not find a requirement connecting a particular AI implementation to retained capability or to the formation of that capability over time. Safety-critical guidance addresses the problem selectively, where failure is acute and immediate.

Two qualifications, and they should travel with that observation. Absence from published instruments is not absence in practice; an organisation may test retained capability under degraded or absent AI support internally. And the review behind this covered accessible mainstream frameworks, not every sectoral or professional obligation in existence.

The World Health Organization's classification of functioning offers a precedent, though not a template. It separates what a person can do in a standardised setting from what they actually do in the setting of their own life, and attributes the difference to environmental conditions.⁶ It is not a measure of unaided performance, and it does not treat support as contamination. What it shows is that capability and situated performance can be held apart, that both can be described, and that a discipline can build a working instrument out of the difference between them.

A caution on use

The second number describes a work system, not a worker. Somebody who cannot function without a system, in a role that was redesigned to require it, has been formed by the design rather than failed by a test. Used as a rating, it becomes one more way of holding individuals answerable for conditions they did not choose. Held the other way round, it is information a worker has as much reason to want as an employer does.

Safety-critical guidance already knows how to frame this where failure is immediate. I did not find in the reviewed corpus an equivalent longitudinal requirement for settings in which judgement may thin slowly and nothing announces itself on the day.

The bill arrives late

Who should be asking is itself a question. Boards sit inside the same incentives that produced the decision. Regulators may have little visibility into how the work is actually done, and professional bodies have interests of their own to protect. What is needed is not externality but standing — somebody with the authority to ask, enough information for the question to mean something, and a horizon long enough that the answer still matters when it arrives. The function implementing a system may not be the one best placed to judge what that system is making harder to see. This is not an accusation of bad faith. It is the ordinary difficulty of assessing an arrangement from inside it.

The lag does much of the work. The bill arrives late, and it arrives to different people. The executive who authorised the implementation has moved on. The consultants have gone. The saving was booked in the year it was made, and the year closed. If the patterns in this paper prove real, they will not present as the consequence of anything. They will present as a vacancy that stays open five months. As the senior person everybody relied on retiring, and the slow realisation

that there is no obvious second. As work that once had three people learning it, now having none. Whether anyone traces those back to a decision taken six years earlier depends on whether anyone was ever asked to.

The mismatch is between the tenure of the people who decide and the time it takes for judgement to mature.

What forms the next person

An organisation might read all of this and conclude that its own interest is reason enough to act. It would be right to act. Its own interest would still not be enough, because the people who could absorb badly redistributed work, argue with it, or put it right were formed somewhere else, over years, by more hands than any single organisation has.

The capability that walks through the door on the first day was made elsewhere. By a previous employer. By a profession. By a training system. By the accumulated practice of a field over a period longer than any employment relationship. An institution can protect its own pipeline scrupulously and still be exposed, because it does not create most of what it draws on.

A sceptical reader will say this is ordinary underinvestment in transferable skill, understood for decades, and that firms have always been reluctant to fund training a competitor might benefit from. That is fair as far as it goes. It does not go far enough here.

Start with what is at stake, which is not principally training. It is practice, judgement, relational knowledge, the accumulated feel for a field that forms while the work is being done — across jobs, across institutions, across a working life. No course delivers that on its own.

Then there is the matter of cost. Formation was never free. Supervision, slower junior output, protected practice, correction, mentoring, the tolerance of mistakes that somebody then had to put right: all of it cost something, and always has. But the cost sat inside the ordinary performance of the work rather than appearing anywhere as an investment in future capability. There is no line item to defend, because there never was one. Reorganise the task and the formative function inside it can go with it, unless somebody deliberately puts something in its place.

This is where the enthusiasm for capturing expertise before it walks out of the door runs into difficulty. The ambition is reasonable and often valuable — what a senior person knows carries on working for the organisation after they have gone. But it answers one half of the problem and may

sharpen the other. If the knowledge is extracted and reused, what forms the next person whose knowledge is worth extracting?

And the consequences do not stay inside the firm. A profession that cannot form its next cohort delivers thinner judgement to every organisation and every client relying on it, including the competitors who also declined to pay.

That combination gives the problem the shape of a commons, which is not a claim that it is a common-pool resource in the classical sense. Knowledge is often non-rival. Capability is instantiated in particular people. Training systems are institutions rather than stocks. What corresponds is the structure: capability formed by many contributors, drawn on across institutional boundaries, vulnerable to local under-contribution, and consequential well beyond whoever makes the immediate decision.

The framing follows the account of agency and its ecologies set out in *The Brain Commons*: capability is lived in persons but never privately produced, and the conditions that make it possible are held in common and can be depleted without anyone deciding to deplete them.⁷ What this paper adds is the point of application — the moment a work system is redesigned, and the decision through which those conditions are altered.

In July 2026, Nolan Lovett reached closely adjacent ground from human resource development. He describes professional expertise as a cognitive commons, regenerated collectively through practice and therefore vulnerable when individually rational adoption decisions remove some of the work through which it renews. His Validation Tether names a related bind: effective AI oversight depends on expertise that AI adoption may itself weaken. His focus is the regeneration of professional expertise across organisations and professions; this paper stays closer to the implementation decision and the conditions through which capability is retained and formed. The two accounts meet around formation. Neither an employer nor a profession can assume the next competent person will simply arrive.¹

Ostrom's lasting contribution was to show that arrangements with this shape are not doomed by it.⁸ Where they hold, it is generally because somebody can see the state of the thing, contributions are visible, and the people who bear the consequences have some hand in the rules. I am not claiming that capability formation has been tested against those conditions, or that it would satisfy them. The correspondence is a place to look.

Formation is easiest to see where somebody was nominally answerable for it. The apprenticeship in the trades. Articles in law. The audit ladder. The second chair in engineering. The registrar years. In each of those, the cost sat inside ordinary work — the slower first draft, the job checked twice,

the hour somebody senior spent not billing. Remove the tasks through which the formation happened and the arrangement is not defunded. It simply stops occurring, and no budget registers the change.

Many workers never enter arrangements like these at all. In platform work, contracted delivery, fragmented employment, small business and much of care, there may be no institution answerable for whether capability is formed, whether workload intensifies, or whether judgement remains possible. Inside an institution there are at least candidate stewards, however weak or conflicted — an employer, a profession, a regulator, a board. A platform may disclaim the employment relationship. There may be no profession to speak of. Regulation may address classification rather than formation. The person carries a risk that has never been assigned to anybody.

For those workers the failure is simply more exposed. Some organisations are drawing down capability that others formed. Elsewhere, people never had reliable access to those conditions to begin with.

Who owns the consequence

Some parts of this problem are already visible to contemporary AI governance. Human oversight, competence and training are increasingly explicit. What is less clear is whether the capability those requirements depend upon is itself treated as something the implementation can alter over time.

The NIST *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* offers one example. It asks that processes for human oversight be defined, assessed and documented; that operator and practitioner proficiency with the system be established alongside them; and that AI-system performance or assurance criteria be demonstrated under conditions similar to deployment.⁹

There is an asymmetry here. The system is to be demonstrated under conditions resembling those it will meet. Human proficiency is defined, assessed and documented, but is not required to be demonstrated under the conditions in which it may later be needed, or connected to how that capability will continue to be formed as the work changes.

The European Union's AI Act reaches further into the human side of oversight. For high-risk systems, Article 14 requires arrangements through which the people assigned oversight can understand the system's capabilities and limitations, monitor its operation, remain alert to automation bias, interpret its outputs, disregard or override them, and intervene or stop the system. Article 26(2) requires deployers to give that role to people with the necessary competence, training, authority and support. Article 4 requires providers and deployers more broadly to support the development of AI literacy among the people operating and using their systems.¹⁰

These are substantive requirements, and the paper would be wrong to suggest that contemporary governance names a human and walks away. It increasingly does not. But the requirements carry a temporal tension. A competent person can be assigned today, trained on this system and authorised to intervene in it. None of that establishes whether the work around them will still be forming the judgement that competent oversight depends on five years from now, or whether the person coming up behind them will encounter comparable conditions through which to acquire it.

At the level of individual AI use, related questions concern cognitive reliance, delegation and integrity. That is a different unit of analysis from the one considered here. That is an important development, but it addresses a different unit of analysis from the one considered here. An interaction can preserve — or impair — cognitive engagement while the wider redesign of work still changes where judgement sits, how expertise is formed and what capability remains available to the organisation over time.

The issue, then, is not that the human is absent from these fields, but that the consequences of a particular implementation can remain distributed across them.

The NHS England announcement is a small illustration of the difficulty. Among the uses it lists for the assistant are drafting clinical letters and registrar training — set out one after the other, as separate benefits. Letter-drafting is one of the ordinary tasks through which a registrar learns what a letter has to carry, and what a consultant sends back. The announcement does not say whether the first was weighed against the second. There would have been no obvious point in the process at which somebody was required to.

Measures relevant to that longer question increasingly exist elsewhere in the same building. Human capital reporting standards now reach across productivity, succession planning, and skills, capabilities and development, with governance and strategy attached to them.¹¹ Workforce planning standards address succession, critical roles and skills development. What I could not find in the reviewed corpus is any requirement to read those measures against a particular AI implementation — when it is approved, or later, once the changed work has had time to shape the people doing it. The succession metric and the implementation decision can still live in different documents, under different functions, without ever being required to meet.

Suppliers enter these frameworks through third-party risk, and the operational technology guidance cited earlier goes further — telling operators to negotiate for transparency about what the AI is doing inside the product, to ask whether it can run without continuous access to a vendor's cloud, and to specify the conditions under which features can be switched off. None of it asks a

supplier to carry the buyer's capability over time. Availability, price, model behaviour, end of life: all settled outside the building, on a commercial horizon shorter than the one that matters here. An institution can own the consequence without owning the conditions.

The question sits across AI assurance, workforce planning, business continuity, work design and professional formation. None of them owns it. What is being settled in the meantime is where judgement will live, how expertise gets formed, and who carries the cognitive load of a faster system.

Capability accountability: six questions before approval

This is the accountability gap the paper is concerned with. An organisation may be accountable for the performance of an AI-enabled system without being required to account for what the implementation is doing to the human capability on which that system will continue to depend.

Capability accountability means making that consequence visible at the point of decision and again over time. It asks what capability the implementation will require, redistribute, form, retain or make harder to recover — and who will remain answerable for that after the immediate business case has been delivered.

Six questions follow. Any substantial AI-enabled redesign should be able to answer them before approval, in the order in which they bear on the decision:

01 What human work—and what part of judgement—is being removed, transformed or newly required?

02 What capability did that work form or maintain, how did it do so, and what will provide those conditions now?

03 Where will displaced cognitive demand and difficult exceptions go, and will the people receiving them have the authority to act?

04 What time is expected to be released, what usable capacity is expected to result, and who will determine what arrives into that space?

05 What capability must remain accessible when the system is degraded, unavailable, outside its competence or wrong?

06 Who owns the consequences for capability and its renewal beyond the horizon of the business case?

None of those requires a new discipline. Each already touches an existing organisational responsibility.

Governance of temporal stewardship is then an emergent quality required to tend to these conditions. A question nobody owns can be answered once, filed, and not asked again. Somebody has to be named to ask, at a moment where the answer can still change the decision, and again later when the consequence has had time to appear. Approval is the obvious first moment. Continuity exercises and annual assurance are the obvious second. Neither needs a new committee, though the implications require a vision of stewardship.

Judgement shock

There is an awkward fact about writing on artificial intelligence now: parts of it begin ageing before they are printed. Systems change. Guidance changes. Standards move through their committees. What looked experimental six months ago becomes ordinary, and concerns that sat at the edge of governance find their way into procurement clauses and assurance schedules.

That movement is visible in this paper's own material. Australian guidance published in 2026 asks organisations to maintain human capability and oversight around AI, to define the competencies required of accountable people, to train those who must intervene, and to preserve alternative pathways when critical systems fail.¹² Government procurement guidance reaches towards internal skills and knowledge transfer rather than leaving capability with the supplier.¹³ ISO/IEC 42105, a standard devoted specifically to human oversight of AI systems, is moving towards publication.¹⁴ These are substantial steps and I would rather they were taken than not. They also begin, for the most part, with a competent human already standing there.

In 1970 Alvin Toffler, working closely with Heidi Toffler, gave the name *future shock* to the disorientation that arrives when change outruns the capacity of people and institutions to absorb it.¹⁵ The technology has changed beyond recognition since. The temporal problem has not, and inside work it has narrowed into something more specific. A work system can be redesigned

in months. The judgement that system still depends on — to supervise it, to recover from it, to know when it is wrong — takes years to form, and forms largely through conditions nobody has ever costed.

I have come to think of that mismatch as *judgement shock*.

It is not the disappearance of expertise, and it is not exhaustion. The word is borrowed less from Toffler than from the clinic, where shock describes a state in which perfusion has already begun to fail while the patient still looks compensated. The observations hold. The person is alert and talking. Nothing announces itself, and then it does. An organisation in judgement shock performs well on the day you look at it. The dashboards are green, the people are competent, the work goes out on time. What has altered is the relation between how quickly the work can be changed and how slowly the judgement it relies on is formed, distributed and renewed.

It shows itself, when it shows at all, in ordinary places. Experienced people who still know exactly what to do, with too much of the operation resting on too few of them. Junior workers meeting fewer of the plain cases through which judgement used to form. A fallback that exists on paper and has not been rehearsed by anybody still employed. A named overseer today, and a thinning path behind them to the next one.

None of that is invisible, and I would not want the idea to rest on invisibility. Each of those conditions can be gone and looked at: how long the vacancy stayed open, when the fallback was last exercised, how the difficult cases are now distributed, who is two years behind the person currently accountable. Judgement shock names a set of conditions somebody could check, and keep checking. If they review and find nothing, that is a result and worth having.

As standards, systems and organisations keep moving, I am less interested in whether each new arrangement can demonstrate a human somewhere in the loop than in the conditions that keep producing humans able to occupy that place. Who gets enough practice to recognise the unusual case. Who learns what a normal failure looks like before being asked to judge an abnormal one. Who has time beside somebody experienced. What must remain possible when the system is degraded, unavailable or wrong? What happens to all of it as the work is made faster, thinner, or differently distributed.

Across law, guidance and standards, the competent human is becoming more visible. What I did not find in the reviewed corpus is a requirement to connect a particular implementation to how that competence will be retained and formed over time.

So this paper will date. I expect the examples to date quickly, and I expect the instruments cited here to be overtaken by better ones. I would be more concerned if the question did.

Work systems can change quickly. The judgement they depend on is formed more slowly. It is exercised by persons and never privately produced, which is why it can be thinned by decisions that nobody experiences as a decision about judgement at all.

The man in the kitchen was not failing to try. Everything he needed was already in the room, and what had altered was the part of him that could receive it. The choices being made now — about what is automated, what remains human, and what the released time is asked to carry — will be answered by people who were not present when they were made, in a decade when the saving is long since booked and whatever bench remains is simply the condition of the work. **Somebody will have to exercise judgement then, with whatever has been left to form it.**

Evidence note

The claims in this paper are of three kinds, set out under "Softer ground" above. This note describes the reading the corpus-level observations rest on.

The observation about retained capability comes from a bounded, non-systematic review of publicly accessible instruments spanning AI risk and governance, regulation and human oversight, safety-critical and operational-technology guidance, business continuity, human-capital and workforce planning, and national or enterprise AI-adoption measurement. The World Health Organization's classification of functioning was read separately, as a precedent for separating capability from situated performance rather than as an instrument of the same kind.

The review was not systematic. It did not cover every sectoral regulator, professional obligation or jurisdiction, and it could not reach unpublished internal practice. Where this paper says something is absent, it means absent from those instruments and nothing further. An organisation that tests retained capability privately would not appear in a reading of this kind. One instrument requires a specific limitation. ISO 30414 was substantially revised in August 2025, and the second edition is now the operative standard. Its published abstract and ISO/TC 260's account of the revision establish what the standard covers, and I have relied on those. I did not obtain the full text. Where this paper observes that something is absent, that observation does not extend to ISO 30414:2025, whose interior I have not read. No absence claim in this paper therefore rests on the interior provisions of ISO 30414:2025.

Notes

1. N. Lovett, "The Tragedy of the Cognitive Commons: How AI Could Disrupt the Regeneration of Professional Expertise", *Human Resource Development Review*, OnlineFirst, 26 July 2026, DOI 10.1177/15344843261470602.
2. NHS England, *500,000 NHS staff to get new artificial intelligence tools to help free up more time for patients*, 8 June 2026. The trial covered more than 30,000 staff across 90 NHS organisations and reported average savings of 43 minutes per person per day, equivalent to about five weeks a year. The assistant is Microsoft 365 Copilot; full rollout is expected by October 2026.
3. On the ironies of automation, including the new demands left to human operators and the erosion of manual competence, see L. Bainbridge, "Ironies of automation", *Automatica* 19(6), 1983, 775–779.
4. CISA and the Australian Signals Directorate's Australian Cyber Security Centre, with the NSA, FBI and partner agencies in Canada, Germany, the Netherlands, New Zealand and the United Kingdom, *Principles for the Secure Integration of Artificial Intelligence in Operational Technology*, 3 December 2025. Skill erosion appears in the risk table and in Principle 1.3; maintaining operational competencies alongside AI in 1.3; failsafe reversion in 2.4; thresholds for defaulting back to non-AI systems in 3.4. All actions taken under it are voluntary.
5. ISO 22301:2019, *Security and resilience — Business continuity management systems — Requirements*, on exercising and testing business continuity arrangements.

6. World Health Organization, *International Classification of Functioning, Disability and Health*, 2001, on the capacity and performance qualifiers and the role of environmental factors.
7. D. J. Norris, *The Brain Commons: Notes on Human Agency and the Ecologies That Shape Us*, May 2026.
8. E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 1990.
9. NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, 2023, at MAP 3.4, MAP 3.5 and MEASURE 2.3. The framework is voluntary guidance rather than a regulatory requirement.
10. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 4 (AI literacy), Article 14 (human oversight of high-risk AI systems) and Article 26(2) (competence, training, authority and support of persons assigned human oversight). Consolidated text as at 27 July 2026 consulted.
11. ISO 30414:2025, *Human resource management — Requirements and recommendations for human capital reporting and disclosure*, 2nd edn, published August 2025, superseding ISO 30414:2018 (withdrawn). The second edition introduces auditable requirements alongside recommendations, establishes a baseline of fourteen required metrics applicable to organisations of any size or sector, and retains eleven human capital areas including productivity, mobility and succession planning, and skills, capabilities and development. ISO/TC 260 records that the revision strengthens guidance on data governance, digital tagging, artificial intelligence and workforce upskilling. This paper relies on the published abstract and on ISO/TC 260's account of the revision; the full normative text was not obtained, and no claim is made here about its interior provisions. Also ISO 30409:2016, *Human resource management — Workforce planning*. ISO/IEC 42001:2023 was considered relevant to the governance landscape but was not included in the bounded absence review because its full normative text was not reviewed for this paper.
12. National AI Centre, *Guidance for AI adoption: implementation guidance*, Australian Government, 2026.
13. Australian Government Department of Finance, *National Framework for the Assurance of Artificial Intelligence in Government — Cornerstones of assurance*, procurement guidance, 2026.
14. ISO/IEC FDIS 42105, *Information technology — Artificial intelligence — Guidance for human oversight of AI systems*, Final Draft International Standard, under approval at the time of writing.
15. A. Toffler., *Future Shock*, Random House, 1970.

Instruments used in the corpus-level review

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023.

Regulation (EU) 2024/1689, Artificial Intelligence Act, particularly Articles 4, 14 and 26(2).

CISA, ASD's ACSC, NSA, FBI and partner agencies, Principles for the Secure Integration of Artificial Intelligence in Operational Technology, 3 December 2025.

ISO 22301:2019, Security and resilience — Business continuity management systems — Requirements.

ISO 30414:2025, Human resource management — Requirements and recommendations for human capital reporting and disclosure (2nd edn, August 2025), reviewed via published abstract and ISO/TC 260 documentation only.

ISO 30409:2016, Human resource management — Workforce planning.

Eurostat: Use of artificial intelligence in enterprises. Statistics Explained, data extracted December 2025, code isoc_eb_ai.

UK Department for Science, Innovation and Technology, *AI Adoption Research*, 2026.

Read separately, and not part of the AI-governance corpus: World Health Organization, International Classification of Functioning, Disability and Health, 2001, cited as a precedent for separating capability from situated performance rather than as an instrument of the same kind.

ABOUT THE AUTHOR

David J. Norris has spent more than twenty-five years in neurological and cognitive practice as a senior occupational therapist, working with people whose judgement, attention and participation had been altered by injury or illness. That clinical ground is what keeps the abstractions in this paper honest: capability is not a property of a person alone, but of the conditions around them.

He writes on human capability, brain health, technology and institutional conditions, and publishes *Notes on Human Agency*. *The Brain Commons: Notes on Human Agency and the Ecologies That Shape Us* is his inquiry into the shared conditions through which agency and participation become possible.

Agency is lived in persons, but never privately produced.

NOTES ON HUMAN AGENCY

Continue the conversation.

For questions, reflections, or invitations
connected to The Brain Commons:

David J. Norris

Brisbane, Australia



Read and subscribe

Scan to enter
Notes on Human Agency —
the Substack publication

CONNECT VIA LINKEDIN

David Norris
[linkedin.com/in/davidjnorrisau](https://www.linkedin.com/in/davidjnorrisau)