

# The Stewardship of Human Capability

## Leadership in the Age of AI

Dr Paige Williams  
August 2026

The way we navigate today must leave us more capable of navigating tomorrow.

## The argument in brief

Artificial intelligence is changing more than what organisations can accomplish. Decisions about where and how AI enters work also shape the developmental environment of work: what people practise, what they delegate, where authority sits, what they learn and what they become capable of doing next.

This creates two connected leadership responsibilities. *Navigation* is the work of responding intelligently to what is changing around us. *Stewardship* is responsibility for what our way of navigating is developing within people and organisations.

The paper argues that *Stewardship* should not become an attempt to preserve every existing human skill. As human and artificial intelligence become more interdependent, the more foundational concern is Cognitive Independence: the integrated presence of Awareness, Agency and Accountability.

The practical question is therefore not simply where humans should remain in the loop. It is how leaders design human-AI interdependence so that organisations can extend what becomes possible while remaining capable of perceiving beyond what is presented, authoring meaningful movement and learning from consequence. Three questions make that responsibility practical: What are we extending? What must we continue to author? What are we developing?

**The governing proposition:** *The way we navigate today must leave us more capable of navigating tomorrow.*

## The changing shape of capability

There is a great deal we still do not know about how artificial intelligence will change work. We do know that it already is.

AI can analyse information, identify patterns, generate possibilities, produce sophisticated work and make recommendations at a speed and scale that previously required considerable human effort. In some kinds of work, people using AI are already producing better work, more quickly. In others, results are much less reliable, and the boundary is not always obvious.

Alongside performance sits another question: What is happening to us as we become more capable with AI?

For most of the conversation so far, capability has been understood through what the combined human–AI system can accomplish: can we do the work faster, better, at greater scale, or solve problems we could not solve before?

These are reasonable questions. They tell us something important about what AI makes possible, but much less about what happens to human capability in the process.

Every decision about how we use AI changes how work is done. A team that begins with an AI-generated synthesis encounters information differently from one that first makes sense of the material itself. A leader who receives machine-generated options before forming a view begins choosing from a different place. A professional may become capable of work previously beyond them while practising less of what the machine now performs. An organisation that delegates action to increasingly autonomous systems changes where Agency sits and how responsibility may be understood.

None of this makes AI inherently diminishing. Human beings have always extended themselves through tools: writing, calculators and search engines all changed what we needed to remember, know or do for ourselves.

AI extends that history while moving further into activities we have treated as part of cognition itself. It can notice patterns, interpret information, generate possibilities, form recommendations and increasingly act within the processes through which we make sense of the world and decide what to do.

The terrain is changing. So are the tools with which we navigate it.

## Capability is not one thing

Research into human–AI collaboration gives us good reason to resist simple conclusions about what follows.

One of the most useful findings comes from Fabrizio Dell’Acqua and colleagues, who describe a jagged technological frontier. In their field experiment with BCG consultants, participants using GPT-4 on 18 tasks within the technology’s capability frontier completed 12.2% more tasks, worked 25.1% faster and produced higher-quality work. On a complex task designed to sit outside that frontier, participants using AI were 19 percentage points less likely to reach the correct answer.<sup>1</sup>

The frontier is difficult to see in advance. AI can be extraordinarily capable at one task and unreliable at another that appears, from a human perspective, to require similar competence.

I think something similar may be happening to human capability. As we work with AI, the shape of our capability changes: we may become more capable at directing AI, interrogating outputs and integrating what it produces while practising less of some work the machine performs. Some capabilities will deepen, some may diminish and others will emerge because AI makes them possible. In this sense, human capability may become increasingly jagged too.

This is an inference from the changing distribution of work between humans and AI, not a finding of Dell'Acqua and colleagues, and jaggedness should not be read as decline. A person who performs fewer parts of a task unaided may still become considerably more capable overall through working intelligently with AI.

The more useful question is what pattern of capability we are creating, and whether that pattern leaves people and organisations equipped for what they will need to do next.

A large meta-analysis of human–AI experiments reinforces the need for care here. Across 106 experiments reporting 370 effect sizes, human–AI combinations performed better than humans alone on average, but worse than the better of the human or AI working alone. The pattern varied by task: the pooled result was negative for decision tasks, while creation tasks showed a more favourable pattern, although the average synergy effect for creation tasks was not itself statistically significant.<sup>2</sup>

Deciding where AI can be used is therefore only part of the leadership work. We are also deciding how capability will be distributed between humans and machines, under what conditions and with what consequences beyond immediate performance.

Emerging research gives us reason to distinguish performance while using AI from what the human learns or retains through using it. In a randomised writing study, ChatGPT support improved essay performance but did not produce significantly greater knowledge gain or transfer. In a separate survey of knowledge workers, higher confidence in generative AI was associated with less reported critical-thinking effort; the study was observational and does not establish causation.<sup>3 4</sup> AI can reduce cognitive effort, often usefully, but the consequences of that offloading depend on what is offloaded, by whom and under what conditions.<sup>4</sup>

Learning involves effort. Expertise develops partly through encountering information before it has been organised for us, wrestling with problems, making judgements, getting things wrong, experiencing consequences and revising what we understand. When AI changes those experiences, it also changes the developmental environment in which human capability is formed.

There is no reason to preserve effort for effort's sake, nor to preserve a human capability simply because a machine can now perform the associated task better. Some capabilities will become less important, others will change and new ones will emerge.

The leadership question is more discriminating: Which changes in human capability do we need to pay attention to? Underneath it sits a more fundamental question:

**What must remain strong in the human as more of our capability becomes shared with machines?**

## A question of movement

My interest in this question predates generative AI. For some time, my work has been concerned with human movement and behaviour: how people and human systems move from where they are towards what they are trying to create, what enables that movement and what causes it to stall.

ACIA, my original model of Cognitive Independence, describes purposeful human movement through four stages — Awareness, Choice, Intention and Action — and the Gateways and enabling conditions through which movement can proceed.<sup>5</sup> At the foundation sit three capacities: Awareness, Agency and Accountability.

Cognitive Independence is the integrated presence of those capacities: Awareness allows a human system to perceive sufficiently clearly; Agency allows it to author movement; Accountability keeps it connected to consequence sufficiently to learn.<sup>5</sup> ACIA describes how purposeful movement unfolds over time, rather than all human behaviour. It was not developed to explain AI.

AI does give us a new context in which to examine that movement because it can participate in it: influencing what comes into Awareness, shaping the options available within Choice, affecting Intention, participating in Action and mediating the evidence through which we understand what happened afterwards. In ACIA, I call that return Proof: the point at which action becomes evidence that can expand Awareness.<sup>5</sup>

That participation may expand Cognitive Independence by extending what people see, increasing the choices available and making new action possible. It can also narrow what we notice, make one course of action unusually easy to follow, distance us from consequences or remove experiences through which expertise would otherwise have developed. Neither outcome is automatic.

This brings me to two connected leadership responsibilities. *Navigation* is the work of engaging intelligently with what is changing around us. *Stewardship* is responsibility for what our way of navigating is developing within people and organisations.


The two are recursive. The capabilities available to us affect how well we navigate; the choices we make while navigating alter what people practise, where authority sits, what they learn and what they become capable of doing next. Figure 1 shows this relationship and the developmental conditions to which Stewardship attends.

That is the leadership problem examined in this paper: the way we navigate today must leave us more capable of navigating tomorrow.

## STEWARDSHIP / FIGURE 1

**Navigation and Stewardship**

Two connected responsibilities within the same leadership movement

**At the centre: Cognitive Independence****Awareness · Agency · Accountability**

Both responsibilities involve perception becoming action and consequence becoming evidence that can expand Awareness.

**Figure 1. Navigation and Stewardship.** The capability available to us shapes how we navigate. The experience of navigating, in turn, can reshape capability. *Stewardship* attends to the developmental conditions throughout this relationship.

## Navigation

Leadership has always required us to move without knowing everything we would like to know. We gather information, draw on experience and make judgements, knowing that some things will become clear only once we act. Sometimes we can see a direction; at other times, the next move is an experiment through which we begin to discover one.

This is what I mean by Navigation: moving intelligently through conditions we cannot fully know in advance. We notice what is happening, gather information and possibilities, decide where and how to move, and act. Then we pay attention to what happens. The return matters because it gives us information that was unavailable when we made the decision.

Navigation draws directly from ACIA's Awareness, Choice, Intention and Action.<sup>5</sup> Choice develops a genuine field of options; Intention gives that exploration a direction. As action meets reality, what happens can become evidence that expands Awareness. This architecture supports movement when the route cannot be fully determined in advance.

## When the maps move too

Uncertainty itself is not new. What is unusual is that AI is changing both the conditions we are navigating and the means through which we understand them. A leader can now ask a system to identify patterns across information no individual could process, generate options, model consequences or recommend a course of action. AI is becoming part of Navigation itself.

Used well, this expands the field within which leaders navigate. But the quality of the account we receive can make it harder to recognise what remains unknown. A sophisticated analysis may still rest on incomplete information; a plausible recommendation can still be wrong. The system can identify patterns in what is available while remaining blind to what is absent, and its fluency can make a provisional view feel settled.

We therefore have to stay attentive to the changing environment and to the systems through which we are coming to know it. Their capabilities can extend our understanding, while their limitations shape what we may fail to notice. The maps are moving too.

## The next best step

When both the terrain and the maps are changing, waiting for a complete route can become its own form of fragility. Some information becomes available only after we take a sufficiently intelligent step and see what it reveals. The movement itself gives us something further to work with.

This is why I find the question "What is the next best step?" useful. It asks us to bring enough judgement to the situation to move, while remaining open to what the move may teach us. A step matters for what it accomplishes and for what it allows us to learn; that learning can change the direction or the scale of what we do next.

The care required depends on what is at stake. For consequential or irreversible decisions, the threshold for evidence and scrutiny will rightly be higher. Navigation recognises that thought and movement inform one another: some understanding emerges only through action whose consequences we are prepared to encounter.

## Where AI *enters* the movement

AI can increase the information available before we move and help us examine what happens afterwards. It also changes the experience of navigating.

Consider asking an AI system to recommend a course of action. If I form a view and then use AI to challenge it, AI enters my thinking in one way. If I ask AI what I should think before encountering the problem myself, it enters in another.

Neither sequence is universally right. Asking AI first may reduce bias, reveal possibilities and avoid work the machine can perform better. In other circumstances, encountering the problem before seeing the machine's interpretation may matter because the encounter itself develops knowledge, judgement or perspective we will need later.

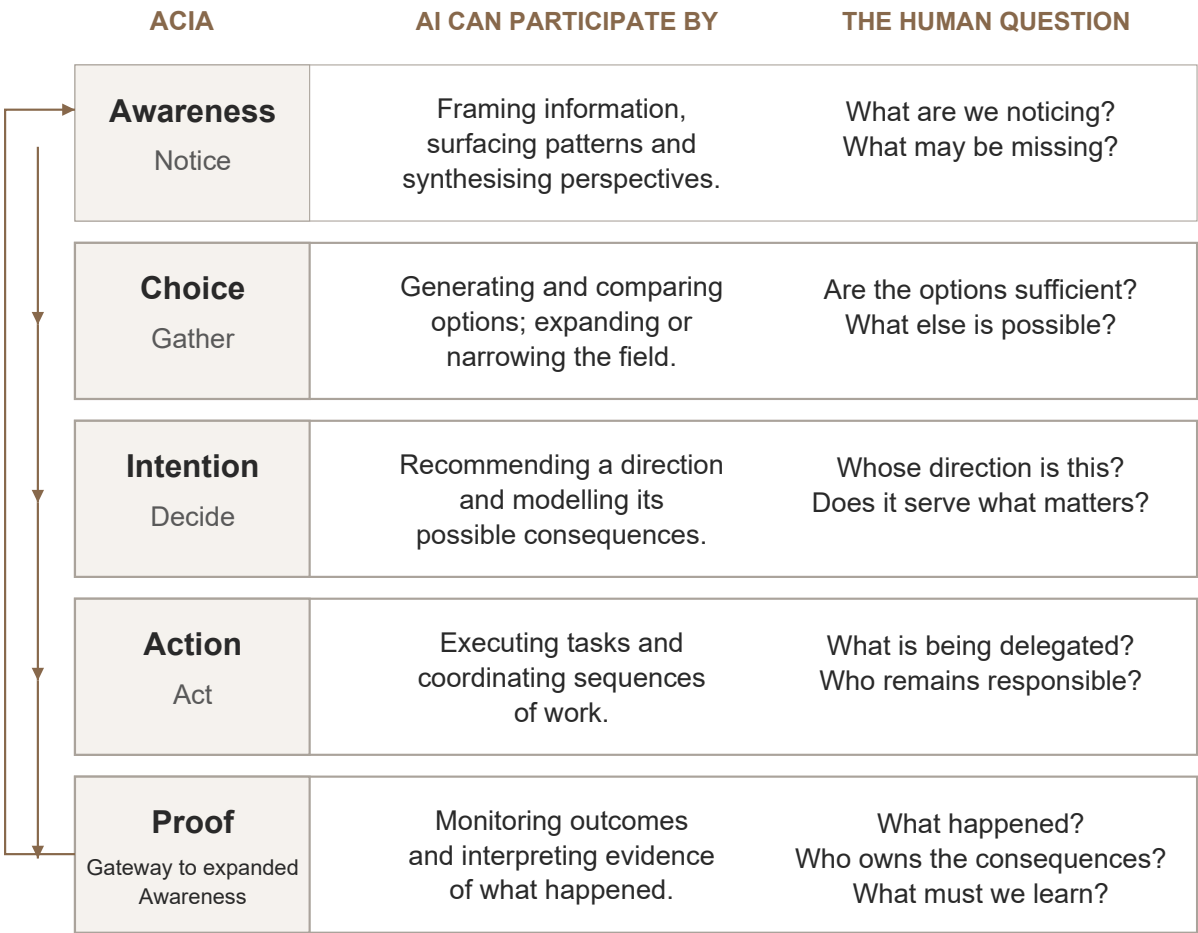
The leadership decision therefore extends beyond whether to use AI to where AI *enters* the movement.

That choice affects what people encounter before information is organised for them, which possibilities they generate themselves, how they participate in decisions, what they experience through action and how they make sense of consequences afterwards. Where AI enters cognition may therefore be as consequential as how much AI we use. Figure 2 locates these questions within the ACIA cycle, showing examples of AI participation alongside the human questions that remain.

STEWARDSHIP / FIGURE 2

Where AI enters purposeful movement

Where AI participates shapes what people encounter, author and learn.



Human involvement depends on the work, the context and what must develop.

ACIA stages: Awareness → Choice → Intention → Action.

Proof: action becomes evidence that can expand Awareness. It is not a fifth stage.

**Figure 2. Where AI enters purposeful movement.** AI can participate throughout the ACIA cycle. The questions help examine where meaningful human involvement matters. Proof is the return through which action becomes evidence that can expand Awareness. Accountability means remaining responsible for consequences and connected to them sufficiently to learn. Adapted from Williams, P. (2026, April), ACIA: A Model of Cognitive Independence; see note 5.

## **Movement produces information**

There are limits to what analysis can tell us before we move. Action creates new conditions: people respond, assumptions encounter reality and consequences appear that were invisible when the decision was made. The world after action contains information the world before action could not.

This is why the return to Awareness through Proof is essential. Proof asks “What happened?”: action becomes evidence that can revise what we understand.<sup>5</sup> It is not confirmation that we were right. That return depends on encountering and integrating the evidence; action alone does not ensure learning.

AI can strengthen this process by detecting signals in outcomes, comparing results with expectations and helping us make sense of what happened. It can also increasingly tell us what happened and what it means.

The issue is not whether humans should make sense of experience unaided. Human sense-making has always been distributed across other people, tools, records, institutions and expertise. The question is whether people and organisations remain sufficiently involved to interrogate the interpretation they receive, recognise when it does not fit the experience and change direction when necessary.

Navigation requires more than access to intelligence. It requires the capacity to remain intelligently involved in the movement.

## **Every move changes the navigator**

When we act, we do more than change the external situation. We practise particular ways of working, exercise some forms of judgement and delegate others, distribute authority and experience consequences. Those experiences can change the organisation, but they do not automatically strengthen its capability.

This connects with a question running through my earlier work on antifragility: why do some people and systems become more capable through challenge while others become more fragile? One idea underneath that work was developmental: disruption can become an experience through which understanding and capability expand rather than simply returning us to a prior state.<sup>6</sup>

AI makes that developmental relationship particularly important. When AI becomes the normal starting point for a problem, we change how people encounter it. When we delegate action, we change where Agency is exercised. When we change human oversight, we change how Accountability is experienced. When synthesis arrives before the underlying material, we may change what people become practised at noticing for themselves.

Some of these changes will make people and organisations more capable; some capabilities may appropriately become less important; others may become newly possible. The leadership question is what kind of capability our choices are developing and whether it leaves them equipped for what comes next.

This is where Navigation alone becomes insufficient. Leadership has to attend both to what is changing around us and to what our way of responding is changing within us. That second responsibility is Stewardship of Human Capability.

## Stewardship of Human Capability

If Navigation is concerned with how we respond to changing conditions, Stewardship asks what our response is developing in us. I use Stewardship of Human Capability to describe leadership responsibility for the conditions in which human capability develops as work becomes increasingly shared with AI.

The object of stewardship is not a fixed set of human skills. Technology has always changed what people need to know and do: capabilities become less important, new ones emerge, expertise changes shape and work moves between people, systems and machines.

The responsibility is to attend to the developmental consequences of how work gets done. Every decision about where AI participates shapes an environment in which people practise some things more and others less, encounter some problems directly and others through machine interpretation, make some decisions and receive others as recommendations, exercise authority in some places and delegate it in others.

Over time, those patterns become developmental. People become practised at what their work repeatedly asks of them.

## Performance and development

The capability of the combined human–AI system and the capability developing within the human are not the same thing. A person may produce substantially better work with AI while becoming less capable of performing some parts of that work without it, and that does not necessarily represent a loss.

If the capability being relinquished is no longer important, the change may be entirely appropriate. The more difficult cases are those in which the activity being offloaded was also doing developmental work.

Making judgements and examining their consequences can develop judgement. Encountering ambiguity can develop the capacity to work with it when the experience is reflected on and integrated. Generating possibilities exercises something different from selecting among possibilities already generated. Taking action and staying connected to its consequences can create opportunities for learning that are difficult to acquire at a distance.

The significance of an activity therefore cannot always be understood from its output alone. Sometimes the work matters partly because of what doing it develops in the person or system doing it.

This concern predates generative AI. In 1983, Lisanne Bainbridge described an irony of automation: operators could retain responsibility for taking over when automation failed while having fewer opportunities to acquire and maintain the knowledge and skills required to do so. She also noted that following computer advice could reduce practice in generating intelligent action for oneself.<sup>8</sup> Her context was industrial process control, but the design problem is recognisable: automation can alter the experience through which people remain capable of effective intervention.

The emerging evidence on AI-assisted work does not justify a general conclusion that using AI makes people less capable. It does justify separating immediate performance from longer-term human learning, cognitive effort and future capability.<sup>3 4</sup> A system designed only around today's output can optimise away experiences that tomorrow's capability depends upon.

## What must remain strong

The case for Stewardship cannot rest on human superiority or on preserving human performance of a task simply because we value the capability associated with it. A machine may make a better recommendation, generate more useful possibilities or execute a decision more reliably.

The more consequential question is: Which human capabilities are foundational enough that leaders should actively steward them even when AI can perform the associated task better?

This is where I return to Awareness, Agency and Accountability. They have a different status from many capabilities we might otherwise put on a list. Judgement, creativity, critical thinking and domain expertise may all matter enormously, but their importance varies with the work and may change as technology changes. Awareness, Agency and Accountability sit underneath those changing requirements: they are conditions through which a human system remains meaningfully involved in its own movement.

Awareness allows us to perceive what is happening, including assumptions, gaps, consequences and possibilities that may not have been presented to us. Agency allows us to author movement: to determine what we are trying to do, choose how to proceed and change direction when necessary. Accountability connects action with consequence strongly enough for us to learn from what follows.

Stewardship does not require keeping AI away from these capacities. It requires attending to whether the way we use AI strengthens or weakens their presence in people and organisations.

## The human system

This is why I use Cognitive Independence as the foundational concern of Stewardship. It does not mean thinking without AI or reproducing unaided everything we can accomplish with machines.

Human intelligence has always been interdependent: we think with other people, cultures, institutions, language, accumulated knowledge and tools. AI extends that interdependence considerably. The question is whether, within it, people and organisations retain Awareness, Agency and Accountability.

Stewardship therefore concerns the human system, not simply the individual worker. By human system, I mean people together with the organisational arrangements through which human perception, judgement, authority and responsibility are exercised. At organisational scale, Awareness is collective; Agency lives in structures and decision rights as well as people; Accountability becomes more complex as outcomes emerge from combinations of human judgement, organisational rules, data, algorithms and increasingly autonomous agents.

## Stewardship is a design responsibility

Seen this way, Stewardship sits inside decisions leaders are already making about how work gets done.

Where AI enters a process, what people encounter before it does, which judgements remain theirs to make, what is delegated, and how people experience the consequences of action all shape the developmental environment of work.

There will rarely be one correct configuration. What makes sense will depend on the work, the person, the level of expertise, the consequences of error, the maturity of the technology and what the organisation needs its people to become capable of doing next.

**Three questions help hold that complexity:**

**What are we extending?**

Where can AI genuinely expand what people and organisations can perceive, understand, create or accomplish?

**What must we continue to author?**

Where does meaningful human involvement remain important, not because humans necessarily perform the task better, but because Awareness, Agency, Accountability, judgement or future capability depend upon genuine participation in the work?

**What are we developing?**

As this way of working becomes normal, what is becoming stronger in the human system? What is becoming less practised? What new dependencies and capabilities are being created? And are those changes ones we are prepared to choose deliberately?

These are questions about designing human–AI interdependence with some understanding of what that interdependence is producing. Figure 3 brings them together as Extend · Author · Develop: connected questions to consider when designing work with AI.

**They also require leaders to hold two time horizons at once:**

**What does this way of working allow us to accomplish now?**

**What is this way of working developing in us for what comes next?**

Navigation attends to how we move through changing conditions.

Stewardship attends to the capability with which we will make the next move.

## STEWARDSHIP / FIGURE 3

**Extend · Author · Develop**

Three questions to consider together when designing work with AI

| <b>EXTEND</b>                                                                                     | <b>AUTHOR</b>                                                                                                      | <b>DEVELOP</b>                                                                                                      |
|---------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b>What are we extending?</b>                                                                     | <b>What must we continue to author?</b>                                                                            | <b>What are we developing?</b>                                                                                      |
| Where can AI expand what people and organisations can perceive, understand, create or accomplish? | Where does meaningful human direction and participation matter for judgement, responsibility or future capability? | As this way of working becomes normal, what strengthens, what becomes less practised, and what dependencies emerge? |
| <hr/> New possibilities<br>Greater reach<br>Expanded capability                                   | <hr/> Purpose and direction<br>Meaningful decision rights<br>Authority to challenge                                | <hr/> Practice and learning<br>Connection to consequence<br>Capability for what comes next                          |

---

**Hold two time horizons together**

What can we accomplish now? What capability are we developing for what comes next?

**Figure 3. Extend · Author · Develop.** These are connected questions, not sequential steps. Together they hold immediate performance and the capability developing for what comes next.

## Cognitive Independence

Cognitive Independence rests on the integrated presence of Awareness, Agency and Accountability. Each does different work, and each is affected differently as AI enters human movement.

### Awareness

Awareness is our capacity to perceive what is happening sufficiently clearly to move intelligently within it.

AI can extend that capacity dramatically: searching beyond the limits of individual attention, detecting patterns across enormous datasets, surfacing perspectives we have not considered and making accessible knowledge we did not previously possess.

But more information is not the same as greater Awareness.

Every system of sense-making selects. AI models generate representations of a situation rather than the situation itself. What they surface can expand our field of view; what they omit, privilege or frame can shape it too. This matters especially when AI moves from being one input into sense-making to becoming the place where sense-making routinely begins. The Stewardship question is whether people and organisations retain the capacity to notice beyond the frame they have been given.

At organisational level, Awareness is distributed across people, conversations, data, relationships, systems and perspectives. AI can connect and interpret that distributed intelligence at a scale previously unavailable. It may also create new forms of convergence.

Research in one bounded creative context provides a useful example. In an experiment involving short-story writing, access to generative-AI ideas increased ratings of novelty and usefulness, while AI-assisted stories were more similar to one another than stories written without AI assistance. The study concerns creative writing, not organisational sense-making, so it does not establish that AI narrows collective organisational cognition. It does, however, make the following question worth testing:<sup>7</sup>

If many people begin sense-making from similar models, summaries and recommendations, could individually expanded Awareness coexist with collectively narrowed Awareness?

The capacity to see more is valuable. So is the capacity to see differently.

### Agency

If Awareness concerns what we can perceive, Agency concerns our capacity to move. It involves authorship: determining what we are trying to do, choosing among possibilities, committing to a direction and altering that direction when circumstances require it.

AI can increase Agency substantially. It can give an individual access to expertise they could not afford, capabilities they have not developed and forms of action that would previously have required a team.

Yet capability and authorship are not identical. Imagine a leader presented with three courses of action generated by AI who reviews them and selects one. A choice has been made. The question is how much of the movement they have authored.

There is no universal answer. The important question is whether the person remains sufficiently involved to understand what is being chosen, determine whether the available options serve what is actually important, reject the frame when necessary and generate another direction when none of the options is adequate.

This is why approval and authorship are different acts. A human can remain formally in a decision process while exercising very little meaningful Agency; conversely, a person can use AI extensively to explore, challenge, analyse and execute while remaining unmistakably the author of the movement. The distinction lies less in how much work the machine performs than in where meaningful direction resides.

At organisational scale, Agency lives in structures as well as people. Decision rights, incentives, authority, technology and workflow determine what options are actually available and who can author a direction. As AI systems move from generating recommendations towards taking action, decision-making and execution are increasingly distributed across humans and machines. In ACIA, Agency remains the human capacity to author that movement.

The governance question is who is authorised to decide. The deeper leadership question is whether the organisation remains capable of determining what is worth deciding towards.

## Accountability

Movement produces consequences. Accountability reconnects us to them.

Within ACIA, Accountability is more than being held responsible for an outcome. It is the capacity to stay connected to consequence sufficiently to learn from it.<sup>5</sup> Candour supports passage through Proof by helping us encounter what happened without defending or rewriting it.

We choose. We act. Something happens. We encounter the difference between what we expected and what occurred. Accountability keeps us in relationship with that information so that consequence can become learning, expanding the Awareness with which the next movement begins.

AI complicates this relationship because decisions can be distributed across people, models, datasets, organisational rules and autonomous agents. Causality and responsibility can separate.

Governance frameworks rightly address formal accountability: who is responsible, who has oversight, what can be audited and where human authority must remain.

For Stewardship, there is another question: Who is close enough to the consequence to learn?

An organisation can allocate responsibility perfectly on paper while creating work in which people experience very little relationship between their choices and what follows. AI can also make consequences more visible, trace decisions more accurately and provide richer evidence from which to understand what happened. The concern is whether people remain sufficiently connected to the feedback loop to learn from what follows.

## Why the three belong together

Awareness without Agency leaves us able to see without being able to move. Agency without Awareness gives us movement without sufficient orientation. Awareness and Agency without Accountability allow us to perceive and act without necessarily learning from what follows.

Together they create a recursive capacity: we perceive, author movement, encounter consequence, learn and perceive again. This is Cognitive Independence.

AI may participate in every part of that movement. Cognitive Independence does not require otherwise.

The question is whether the human remains sufficiently present to perceive beyond what is presented, author what is worth pursuing and take responsibility for what the combined system brings about. That is why Awareness, Agency and Accountability are foundational rather than simply desirable capabilities: they do not prescribe which tasks humans should retain; they describe conditions through which humans can remain active participants in increasingly shared intelligence.

### **Cognitive Independence within deep interdependence**

There will be things people and organisations can accomplish with AI that they simply cannot accomplish without it. Requiring people to retain unaided competence in everything AI helps them do would confuse capability with self-sufficiency.

The aim is to participate in an increasingly interdependent system without surrendering the human capacities through which direction, responsibility and learning remain possible. This is Cognitive Independence within deep interdependence.

The question therefore shifts from what this person or organisation can do to what they can still see, author, own and learn from as more of what they do becomes shared with AI. That is the human capability Stewardship is ultimately concerned with.

## Navigation + Stewardship

Navigation and Stewardship describe two connected responsibilities within the same leadership movement.

Navigation attends to changing conditions and how we respond: we notice, gather options, author a direction and act. In ACIA, perception becoming action is the inside-out direction of purposeful movement.

Stewardship attends to what that way of responding is developing in the human system. As consequence returns through Proof to reshape Awareness, the movement is outside-in. Both responsibilities involve both directions: capability shapes how we navigate, and the experience of navigating can reshape capability when it is encountered and integrated.

At the centre of both sits Cognitive Independence. The movement does not end when we act: what happens next changes both the external conditions and the people and systems operating within them, and then we begin again.

### Two responsibilities, one movement

Leadership necessarily attends to what is changing around us. There is no serious version of Stewardship that asks an organisation to develop human capability without regard for its capacity to respond successfully to its environment. Navigation remains essential.

But what begins outside eventually becomes part of the inside. Changes in technology become changes in work; changes in work alter what people practise, delegate and learn, where authority sits and what they become capable of doing. The reverse is equally true: what exists inside the organisation shapes what it can perceive, choose and do in the world.

The relationship is recursive: The world changes. We navigate. Navigation changes us. What changes in us affects how we navigate next.

AI participates throughout that movement. It can extend what the system is capable of at each point while also altering what people become practised at doing for themselves. Leadership has to work with both realities at once.

### Capability has a history

This changes the time horizon of an AI decision. Consider two organisations adopting the same technology to perform the same piece of work. Both may achieve a similar immediate performance improvement, but they may not be building the same organisation.

In one, AI could progressively distance people from the underlying problem, the judgement involved and the consequences of decisions being made. In the other, it could remove work that no longer deserves human attention while creating room for deeper judgement, better questions, experimentation and more consequential human contribution.

Real organisations will contain mixtures of both. The point is that the performance metric alone cannot tell us what capability is accumulating. Capability has a history.

Today's ways of working become tomorrow's experience base. Repeated choices become habits; habits influence expectations; expectations affect what people notice, what they attempt, what they

delegate and what they believe is theirs to own. Over time, the way an organisation uses AI becomes part of what that organisation is capable of becoming.

## **More capable for what?**

There is an important difficulty in the proposition that the way we navigate today should leave us more capable of navigating tomorrow: more capable for what? We cannot know precisely what future capability will be required. If we could, Navigation would be a much simpler task.

Nor should leaders preserve every capability on the possibility that it may someday become useful again. Stewardship is not an argument against specialisation, delegation or technological progress. The standard has to be more foundational.

We need enough Cognitive Independence to encounter what we do not yet know; enough Awareness to perceive change rather than simply receive an account of it; enough Agency to choose a direction when the available options are inadequate; enough Accountability to remain connected to consequence, learn and update.

These capacities do not tell us what tomorrow will require. They preserve our ability to find out.

That is the sense in which more capable matters here. It does not mean more capable of reproducing yesterday's work without AI. It means more capable of participating intelligently in whatever human-AI system comes next.

## **No permanent boundary**

The relationship cannot be designed once and left in place. Human capability changes. AI capability changes. The work changes. A configuration that extends capability today may become constraining later; a capability worth exercising directly while expertise is developing may reasonably be delegated once that expertise exists.

There is no permanent boundary to establish between human and machine. The leadership responsibility is to keep shaping the relationship deliberately as both sides of it change.

## Remaining capable

AI capability will continue to change, and so will the boundary between what humans do and what machines do. Stewardship should not become an attempt to preserve human capability as it currently exists. Some capabilities will become less important. Others will emerge. Work that machines can do better should not remain human simply because it once was.

The responsibility is to shape the changing relationship deliberately, taking seriously both what AI makes possible and the developmental conditions our choices create.

### Three questions keep that responsibility practical:

What are we extending?

What must we continue to author?

What are we developing?

Together, these questions hold two time horizons: what this way of working allows us to accomplish now, and what it is developing in us for what comes next.

The central question of this paper has been: What must remain strong in the human as more of our capability becomes shared with machines?

My answer is Awareness, Agency and Accountability. Not because these are the last things AI will be unable to do; AI already participates in the processes through which they are exercised. They matter because together they allow people and organisations to perceive what is happening, author movement and remain connected to consequence strongly enough to learn.

That is Cognitive Independence within deep interdependence.

We cannot know precisely what capabilities the future will require, and we cannot preserve our way into readiness for it. What we can steward are the conditions that allow people and organisations to encounter what they do not yet know, choose within it, act and learn.

Navigation keeps us responsive to what is changing around us. Stewardship keeps us attentive to what our response is developing within us.

The task is to hold both: to extend what becomes possible, remain authors where authorship matters, and pay attention to what our choices are developing for what comes next.

We cannot know the whole route from here. We can remain capable of finding our way.

The way we navigate today must leave us more capable of navigating tomorrow.

## Notes and references

1. Dell’Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2026). “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality.” *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>.
2. Vaccaro, M., Almaatouq, A., & Malone, T. W. (2024). “When combinations of humans and AI are useful: A systematic review and meta-analysis.” *Nature Human Behaviour*, 8, 2293–2303. <https://doi.org/10.1038/s41562-024-02024-1>.
3. Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). “Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance.” *British Journal of Educational Technology*, 56, 489–530. <https://doi.org/10.1111/bjet.13544>.
4. Lee, H.-P. (H.), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). “The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers.” *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, Article 1121, 1–22. <https://doi.org/10.1145/3706598.3713778>. The study is observational and based on self-report; it does not establish causal skill loss.
5. Williams, P. (2026, April). ACIA: A Model of Cognitive Independence. Dr Paige Williams. ACIA is Williams’ original conceptual model of purposeful human movement. It distinguishes the foundational capacities of Awareness, Agency and Accountability from the stages Awareness, Choice, Intention and Action. Its Gateways are Possibility, Permission, Pathways and Proof; the associated Unlocks are Curiosity, Courage, Commitment and Candour. Unlocks are enabling conditions, not guarantees of passage or substitutes for authority, resources or other contextual conditions. Throughout this paper, Cognitive Independence is used in this ACIA sense.
6. Williams, P. (2024). *The Leader’s Ecosystem: A Guide to Leading, Exceptionally*. Grammar Factory Publishing. The antifragility discussion describes becoming antifragile as an active process of expanding understanding and capability through uncertainty, randomness and disruption, and asks how we can be better coming out of an experience than going in.
7. Doshi, A. R., & Hauser, O. P. (2024). “Generative AI enhances individual creativity but reduces the collective diversity of novel content.” *Science Advances*, 10(28), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>.
8. Bainbridge, L. (1983). “Ironies of automation.” *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8). Bainbridge’s analysis concerned industrial process-control automation and the problems created when operators retain responsibility for abnormal conditions while opportunities to acquire and maintain relevant manual and cognitive skills are reduced.

## About Dr Paige Williams

Dr Paige Williams works with CEOs, executive teams and boards on the human and leadership challenges that arise when the way forward is uncertain and the stakes are high. A PhD in Organisational Behaviour informs her work, alongside original thinking and decades of experience with leaders and organisations. Her work on accountability, antifragility and human capability examines how people and systems can remain capable of intelligent movement through complexity, disruption and change.

Her current work on Leadership in the Age of AI explores what must remain strong in the human as intelligence and capability become increasingly shared between people and machines.

## Continue the conversation

The Stewardship of Human Capability can be explored with CEOs, executive teams and boards as a keynote, executive briefing or facilitated leadership conversation.

[drpaige.au](http://drpaige.au)

## Suggested citation

Williams, P. (2026, August). The Stewardship of Human Capability: Leadership in the Age of AI. Dr Paige Williams.