

ADOPT

AI & HUMAN BEHAVIOUR





Adopt

If AI tools are to deliver on their promises, from increasing productivity to reducing administrative burdens, they need to be widely and meaningfully used. However, adoption is not a switch you can flick on. It is a behavioural process shaped by habits, heuristics, emotions and social context. Behavioural science can help understand this process and design for it.

Much of the discourse around AI treats adoption as a yes/no question: do people adopt or don't they? However, this framing is too simplistic. Adoption is not binary: It runs along a continuum from shallow to deep adoption.

Moreover, adoption plays out in different contexts: individual vs organisational, consumer vs professional, public vs private sector. The barriers and enablers differ according to these contexts.

Here, we focus on the adoption of AI within professional firms. This includes desk-based workers, as well as clinicians, public servants and other professionals. These individuals are likely to be exposed to AI during their work, but the way they use AI is shaped by organisational norms and leadership.

■ Where AI adoption is really happening

Often, the most valuable uses of generative AI don't come from formal deployments. According to an April 2025 HBR <u>analysis</u>, top GenAI use cases include therapy, personal organisation, learning, creative projects, and self-reflection. Inside organisations, adoption may be happening in the shadows as individuals use their own AI tools to perform tasks. These use cases might be shaping adoption far more than we anticipate.



Recent MIT research on over 300 Al initiatives finds that while over 40% of organisations have piloted general-purpose LLMs, only 5% have implemented embedded or task-specific Gen Al. The authors term this the 'GenAl Divide'.

▲ A behavioural framework for AI adoption

To understand and improve adoption for professional workers, we propose a behavioural framework built around three stages: no adoption, shallow use and deep use.

Of course, organisations can operate at varying levels within this framework for different tasks, and individuals within an organisation will also be at various stages.

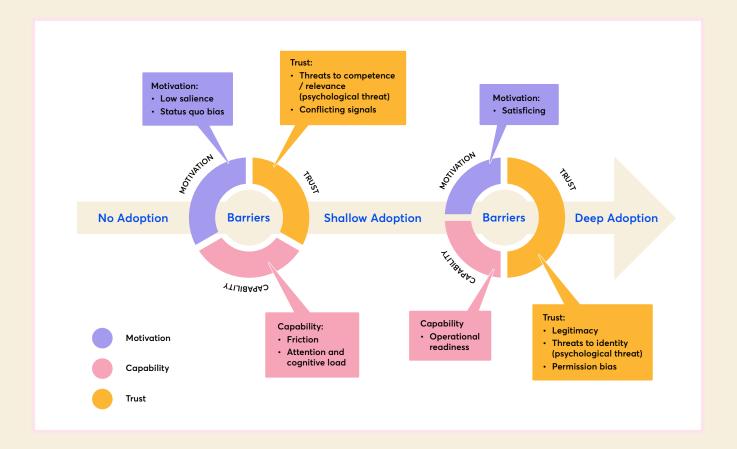
Shallow vs deep adoption				
	Shallow	Deep		
Tasks	Simple, low-stakes tasks such as writing emails or summarising text	Applied to complex, high-value or core tasks, eg, decision-making		
Usage	Ad hoc use for narrow tasks	Embedded use across processes		
Frequency	Occasional or reactive	Frequent and proactive		
Integration	Standalone tools	Integration into workflows		
Impact	Marginal productivity gains	Strategic gains in quality, efficiency and/or innovation		

If no adoption, shallow adoption and deep adoption are specific points along the continuum, then motivation, capability, and trust are the drivers of movement along the continuum. This model is informed by established behavioural science frameworks like COM-B. While crucial environmental factors like choice architecture and friction are often categorised under 'opportunity', we address them within capability, as they directly impact whether users feel able and confident to integrate Al into their workflows.

- → Motivation: whether people see a clear, desirable reason to use Al
- → Capability: whether they feel able to use it effectively and confidently
- → Trust: whether they believe the AI aligns with their values

Each of these factors has its own behavioural underpinnings and solutions.





No -> Shallow Adoption		
	Barrier	Enabler
Motivation	Low salience	create implementation intentions; frame messages to staff; use messenger effects; harness social norms; foster trust through operational transparency
	Status quo bias	use commitment devices; draw on behavioural design
Capability	Friction	harness choice architecture; run sludge audits
	Attention and cognitive load	replace existing work rather than add to it; encourage experimentation; create AI champions
Trust	Threats to competence / relevance (psychological threat)	increase exposure; highlight unique human expertise; frame messages to staff; personalise the experience for staff
	Conflicting signals	provide incentives; establish a clear mandate and guardrails; showcase success to build momentum

	Shallow -> Deep Adoption	
	Barrier	Enabler
Motivation	Satisficing	create implementation intentions; frame messages to staff; use messenger effects; harness social norms; foster trust through operational transparency
Capability	Operational readiness	signal institutional support; encourage bottom-up adoption rather than top-down; structure the adoption journey
Trust	Legitimacy	increase exposure; avoid AI exceptionalism in framing; anthropomorphise AI; embed transparency; evaluate impact and embrace the results (positive or negative
	Threats to identity (psychological threat)	harness loss aversion; democratise AI adoption; use social proof
	Permission bias	signal clearly; use sandboxes

We have structured the adoption challenge around both individual-level and organisational-level barriers, recognising that some barriers to adoption rest outside the individual's control. Similarly, we propose behavioural enablers that can be leveraged at both the individual and organisational level.

This framework is not intended to be a technical taxonomy. Rather, it focuses on what people actually do. This model allows us to better understand where people are on their adoption journey and what can be done to improve it.

▲ Technical taxonomies

Long before behavioural science became the established field it is today, Everett Rogers' Diffusion of Innovation Theory set out five key factors that influence how new ideas and technologies spread.

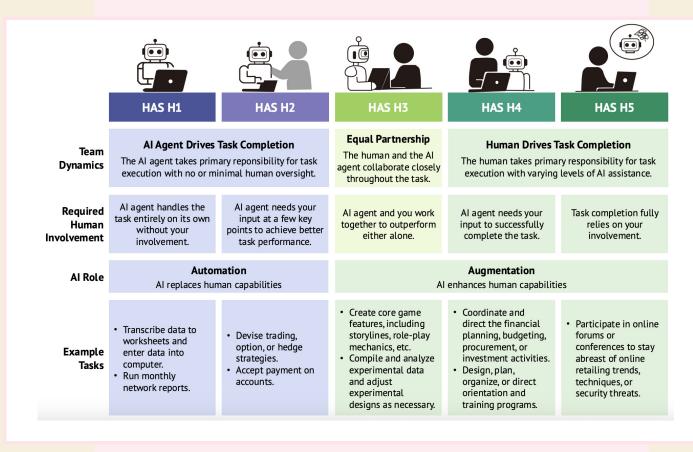
- 1. **Relative advantage**: people are more likely to adopt Al if they perceive it as better than what it replaces. But this advantage must be both visible and specific.
- 2. Compatibility: Al needs to align with a person's values and norms.
- 3. **Complexity**: All needs to be easy to use. Generative Al, of course, has extremely high usability.



- 4. **Trialability**: Al feels safer when people can try it first, particularly in lower-stakes
- 5. **Observability**: if users can see the benefits of Al, they are more likely to adopt it.

Each of these is a fundamentally behavioural question and remains central to understanding the key drivers behind adoption.

Alongside Rogers' classical theory, contemporary frameworks are emerging to characterise Al use. **Researchers at Stanford** have created a Human Agency Scale (HAS) and provide five levels of Al roles:



This taxonomy is useful for mapping automation risk and system design. However, we take a different lens, focusing on human behaviour. The relevant distinction is not between automation vs augmentation, but whether individuals and organisations can discern when and how to choose between them



✓ From no adoption to shallow use

The first step on the ladder is to use any Al tool. However, behavioural factors mean that even this can be surprisingly difficult.

Motivation

Low salience

Professional workers are highly aware of Al. In McKinsey's <u>Al in the</u> <u>workplace report</u>, they find that 94% of those surveyed report having some familiarity with generative Al tools. Awareness does not seem to be a barrier for this group.

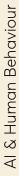
However, the salience of Al's perceived benefit is low, due to:

- → **Benefits being abstract**. For example, much of the discourse around Al refers to 'increased productivity' or 'improved decision-making'. These terms can be disconnected from individuals' specific tasks or issues.
- → **Hidden wins**. Al tools can speed up work or reduce effort, but it's not always possible to quantify these gains. Al may also help in small ways across a variety of tasks. These benefits may not accumulate into a memorable sense of overarching impact.
- → **Automatic processing**. For shallow tasks, Al's assistance may be so seamless that it goes unnoticed. Because the help requires little conscious thought, users don't mentally register the benefit, leading them to undervalue Al's cumulative impact.

▲ Increasing salience of the benefits of AI

What can individuals do?

 Create implementation intentions. Creating simple plans to use AI may yield higher adoption. Simple 'if-then' type plans ('implementation intentions') may be particularly effective. They break goals down into specific actions, reducing the cognitive load of decision-making. The evidence shows this approach has been effective across a wide range of behaviours.





№ What can leaders do?

- 2. Frame messages to staff. Translate abstract benefits into concrete task-linked outcomes, eg, rather than saying Al 'boosts productivity', say Al 'helps write drafts 40% faster'. For example, BIT research for the UK Department for Science, Innovation and Technology found that Al-assisted literature reviews were completed in 23% less time. Making the benefit concrete may encourage other researchers and firms to adopt a similar approach.
- 3. **Use messenger effects.** The choice of messenger also affects adoption. In South Korea, <u>researchers</u> found that messages from supervisors encouraging adoption had a positive effect on intention to adopt. The generalisability of the messenger effect may, of course, vary depending on national cultural dimensions such as <u>power-distance</u>.
- 4. Harness social norms. Highlight that others have used Al and benefited from it, which may increase adoption. In one study, participants were shown that others had successfully used an Al chatbot without issues. Seeing a social norms message increased perception of Al's personal and social benefits. Another study showed participants were far more willing to use an algorithm when making stock price forecasts when told that a majority of other users also used the algorithm. A pilot study in hospital emergency departments incorporated nudges into an Al-powered clinical decision support tool, resulting in an increase in the adoption rate from 21% to 39%. The nudges included displaying peer comparison data.
- 5. Foster trust through operational transparency. In a basic sense, operational transparency involves being explicit about when AI is being used. But it also involves showing how AI is working and what it is doing. Many LLMs now provide users with their reasoning processes. This provision can trigger the 'labour illusion effect,' whereby users trust and value results more when they see effort has been invested in their creation (ie, users can see that the LLM has exerted effort to produce its results).



Status quo bias

Humans have a tendency to <u>prefer the current state of things</u>, even when a better alternative may be available. So people may not use AI even if they are aware that it means they can do things quicker or with less effort.

In an <u>experiment</u> we ran at BIT, we showed participants UK government webpages on different topics (from private renting to constipation in children) and asked them to find answers to specific questions relating to these pages. Participants were randomly allocated to five treatment conditions, with four-fifths of the participants seeing a chatbot and one-fifth seeing no chatbot. We found that only 40% of those seeing a chatbot chose to message it. While a significant proportion may have been able to find the information through traditional webpage navigation, this finding still highlights that people might stick to using what they know.

Overcoming status quo bias

What can individuals do?

 Use commitment devices. Individuals can decide in advance to increase the costs they incur by not using AI. For example, individuals can publicly declare their intention to integrate AI into specific tasks (eg, using AI to take meeting notes). If they feel that failure will mean a hit to their reputation, they will be more motivated to follow through.

What can leaders do?

2. **Highlight tipping points.** Evidence from Microsoft's analysis of 1,300 Copilot users illustrates how small gains can become behaviourally meaningful when they accumulate. They find that just 11 minutes of daily time savings is enough to act as a tipping point where users begin to perceive the tool as valuable. After 11 weeks of consistent use, the majority reported that Copilot had fundamentally improved their productivity, enjoyment of work, and work-life balance. This '11-by-11 rule' demonstrates how minor, often unnoticed wins can compound into habit formation, which in turn reshapes workplace norms.



3. **Draw on behavioural design**. In the BIT <u>ChatGOV experiment</u> mentioned above, 53% chose to message the chatbot when it was shown on a whole page. That figure was significantly higher than when the chatbot was shown in a smaller box in the corner of the page (31% and 26% engagement rates). Leaders can work with their technology teams to draw on these lessons and design tools and interfaces to maximise engagement.

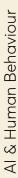
Capability

Friction

Chat-based generative AI is generally highly usable, particularly for lower complexity tasks. But adoption can stall if it is not integrated into existing workflows and systems: workers may feel overloaded with rules determining appropriate use, privacy concerns, and how to expense AI tools.

Reducing friction

- **№** What can leaders do?
- 1. Harness choice architecture.
 - a. **Set defaults**. Organisations can explore automatic integration of AI where suitable. For example, this could be enabling an AI notetaker across an organisation that starts automatically when an internal meeting starts. However, defaults are powerful and should be used with care: unthinking use of defaults may risk ethical or privacy concerns.
 - b. Reduce effort. Even if an AI tool is not made the default, integrating it within workflows means that users can access it with less effort and that makes a difference. For example, embedding an AI chatbot on an intranet site can help users find information about key policies.



- a. **Create timely prompts**. Prompts can make Al tools salient when they are needed. For example, a CRM system could present an Al-generated summary of recent client interactions when a user opens a client's profile, making relevant insights immediately available. **Researchers** found that providing just-in-time, contextually relevant information within a software development environment led to increased task efficiency and user satisfaction compared to static help documents.
- 2. **Run 'sludge audits'**. For large organisations, leaders may wish to systematically identify and reduce points of friction in the adoption of AI a 'sludge audit'. Researchers have had preliminary success with using AI itself to detect patterns of friction that can be harmful.

Attention and cognitive overload

Al may not be given sufficient attention at an organisational level, meaning it sits in the pile of 'nice-to-dos' rather than 'must-dos'. Leaders may be focused on the present and fail to divert enough mental effort to Al adoption.

Addressing low attention and cognitive overload

What can leaders do?

- 1. **Replace existing work rather than add to it**. Reduce cognitive load by framing AI as a way of substituting tasks, rather than adding to them. Leaders can support staff to map current workflows and identify routine or time-consuming tasks that AI might help automate. They can also update performance metrics so that measures reflect the new AI-enabled process. They must then signal to the organisation that AI use is expected for certain tasks.
- 2. **Encourage experimentation**. Small, deliberate trials (eg, a <u>test & learn approach</u>) can help kickstart Al adoption without requiring a comprehensive, resource-intensive plan upfront. By launching focused pilots, leaders can gather quick, real-time feedback on Al's utility. Early wins from these trials can create positive attention, build buy-in and provide evidence to support scaling adoption



- more broadly. An iterative approach allows organisations to refine tools 'in flight' rather than delaying adoption until every detail has been perfected.
- 3. **Create AI champions**. Leaders should not be the only people encouraging adoption: AI champions can enable learning across teams. **Research** on innovation champions shows that champions are most effective when they are: early adopters with a personal interest in the technology; well-connected across the organisation through informal networks; and demonstrate credibility and enthusiasm. Champions should be supported by leaders by giving them: dedicated time to explore tools and help others; priority access to training; organisation-wide recognition; platforms to share their findings.

Trust

Threats to competence/relevance (psychological threat)

Status quo bias is passive: a preference to avoid change. Psychological threat, in contrast, can create a more active resistance to adoption. Threat may take the form of:

- → Perceived self-salience. A recent set of <u>experiments</u> found that individuals who have a clear picture of their future career ('future work self salience') feel more in control of their career trajectory after interacting with Al. Those with low future self-salience saw a reduction in their sense of agency after interacting with Al. This suggests that for some individuals, especially those uncertain about their future role or value in the workplace, Al may heighten existential concerns rather than offer empowerment.
- → The moralisation of AI. There is evidence to suggest that AI is becoming moralised among some individuals, ie, subject to 'conversion of an object or activity preference into something with negative moral status'. Across two studies, researchers found that opposition to AI accounted for a minority of participants (11-39%). For those opposing, the objections were of a moral nature in three out of four applications of AI: AI-generated art, AI companionship and legal AI. Opponents indicated that their views would remain unchanged even if the AI applications were proven to be beneficial.



→ **Psychological reactance**. Individuals may resist when they feel that Al has been imposed on them and they lack freedom of choice. Much of the framing around Al adoption, particularly emphasising 'do or die', is unhelpful.

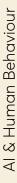
Addressing threats to competence/relevance

What can individuals do?

- 1. Increase exposure. As individuals spend more time using AI, their optimism towards the technology increases. Google and Public First found that AI training led to an increase in positive sentiment towards AI by 22 pp for education workers, 13 pp for small and medium businesses, and 9 pp for trade union members. Similarly, in BIT's ChatGOV trial, mentioned earlier, we found that exposure to AI led to considerably higher support for government use of AI for similar tasks. Evidently, if people feel threatened, they are unlikely to expose themselves in the first place, but small, low-stakes exposures may encourage adoption.
- 2. **Highlight unique human expertise**. Individuals can <u>reaffirm their professional value</u> by focusing on skills that AI cannot replicate. This means emphasising abilities like deep domain knowledge, critical judgment, client relationships, and ethical oversight, which are needed to complement and guide AI-generated outputs.

What can leaders do?

- 3. **Frame messages to staff**. As AI is introduced, it can be framed to mitigate concerns over self-salience and moralisation. Companies could encourage (and incentivise) workers to try tools for just 10 minutes, for a purpose of their own choosing. That can frame AI as a tool for staff to use, rather than something that is imposed on them.
- 4. **Personalise the experience for staff**. Allowing users to personalise their own experience with Al may <u>combat psychological</u> <u>reactance</u>, providing them with more control. This personalisation should be co-designed with workers to maximise its potential effectiveness. Worker consultation is <u>associated</u> with better outcomes for workers.





Conflicting signals

Even where AI tools are available and there is limited friction to their adoption, sending clear signals to staff matters. They can reassure staff that AI adoption is encouraged, even expected. When employees don't receive clear and explicit cues about AI adoption, they are likely to default to existing ways of working.

Unfortunately, organisations often fail to send the right mix of signals to encourage their uptake. Problems include: ambiguous or inconsistent messaging, a lack of visible incentives to adopt, and uncertainties about organisational priorities. Over half of those <u>surveyed</u> across 14 countries (2023) said they were using unapproved generative AI tools because there is a lack of clarity in the company policy.

Creating signals

► What can leaders do?

- 1. **Provide incentives**. Leaders can provide incentives to use Al. This can be a collective incentive for all staff, such as the £1mn bonus pot that <u>one law firm created</u> for its staff if they used generative Al at least 1 million times in a year. Or it can be for individuals, celebrating early Al adopters and producing social influence for others to adopt. But incentives need to be carefully designed to avoid backfire effects and unintended consequences. For example, over-incentivising the quantity of Al use may cause employees to prioritise hitting targets on usage, rather than extracting genuine value from the tools.
- 2. **Establish a clear mandate and guardrails**. Leaders could issue a clear, enthusiastic mandate for Al adoption. This involves not only championing the potential benefits but also providing simple, unambiguous guidelines for approved use. By creating a one-page 'Al Charter' or a simple 'do and don't' list, leaders can remove the fear and ambiguity that causes inaction, giving employees the psychological safety to integrate Al into their work.
- 3. **Showcase success to build momentum**. Go beyond simple recognition by amplifying those who have achieved early Al wins. Successful projects could be turned into internal case studies or demos. Celebrations like these can provide effective incentives for staff.



▲ From shallow to deeper use

Deeper use means embedding AI more fully into tasks, so that it consistently augments or automates. The key is how fully AI is integrated, rather than how intensely it is used. For example, employees frequently using AI to draft emails or summarise uploaded documents represent high intensity. However, integration may be shallow because workflows and approval processes remain unchanged.

Deeper adoption occurs when AI is embedded into the organisation's systems and routines. For example, a customer support team that fully integrates AI into its ticketing platform, allowing AI to triage, prioritise, and auto-resolve common queries, with human agents only handling exceptions. **Krpan et al.** found that even when AI demonstrably improved diversity outcomes and efficiency, professionals were hesitant to move beyond shallow use cases, they might accept AI for CV screening but resist its use in final selection decisions. This **vertical integration** not only increases efficiency but also transforms the underlying workflow, making AI a core part of how work is structured, rather than just a frequently used add-on.

Moving AI adoption from shallow use to deeper use is likely to bring much greater benefits. Only <u>1% of companies consider themselves fully mature</u> in AI deployment, citing organisational barriers and leadership as the key barriers. <u>Accenture</u> finds that those companies with fully modernised, AI-led processes achieve 2.5x higher revenue growth and 2.4x greater productivity than those that haven't used AI. For developers, the use of AI yields benefits beyond marginal gains. An <u>RCT with Google software engineers</u> found that integrating AI cut task time by 21%.

Compare that to the shallow use of Al. A recent <u>study in Denmark</u> looked at the adoption of Al tools among 25,000 workers in 11 occupations. The study found no significant impact on wages or hours worked. They estimate that Al saves just 2.8% of work time on average. The researchers posit that one reason for the low impact on productivity could be shallow adoption: "while chatbots may save time on existing tasks, these savings may not increase productivity on marginal tasks unless employers adapt workflows accordingly".

So, what are the behavioural underpinnings of our failure to see deeper adoption of AI? How can behavioural science yield improvements?



What deeper adoption of AI is not: a caution against cognitive offloading

It is easy to see how deeper use of AI might lead to overreliance. Nascent research, often limited by small sample sizes and lacking robust replication, <u>finds</u> a significant negative correlation between frequent AI tool usage and critical thinking abilities. And, as Oliver Hauser and Anil Doshi <u>demonstrate</u>, generative AI has the power to enhance individual creativity, but it can come at the cost of reducing the collective diversity of novel content, ie, it can lead to homogeneity.

Leaders should monitor these risks as adoption deepens, rather than imposing premature restrictions. We explore the wider societal implications of cognitive offloading and strategies for scaffolding human thinking in Adapt.

Motivation

Satisficing

Satisficing, coined by Herbert Simon in 1947, means that people settle for a 'good enough' use case and cease exploring even if more optimal solutions exist. That tactic can provide good outcomes, and some adoption is likely better than none. However, for deeper Al adoption, exploring potential additional, more powerful uses, organisations need to move beyond satisficing.

Going beyond satisficing

№ What can leaders do?

1. Inspire with examples. Staff may simply not know what kinds of deeper use are possible. Leaders must actively demonstrate what's possible. They can curate and share a library of role-specific use cases, showing how AI can tackle complex challenges in finance or marketing, not just simple tasks. For example, the legal team might go beyond using AI to proofread contracts, and instead use it to perform initial discovery: searching and categorising thousands of documents for specific legal concepts. Leaders could even explain their own uses: how they used an LLM as a sparring partner to challenge assumptions in a draft strategy, for example.

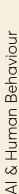
- 2. **Provide incentives**. Run structured short-term competitions to reward experimentation with AI. Wharton professors have developed <u>innovation tournaments</u>, which they claim can move beyond traditional methods such as hackathons. They argue that innovation shouldn't just be the product of a few select employees who are seen as the highly creative types. Instead, they solicit ideas from across an organisation and whittle them down to a few ideas through a structured process.
- 3. **Build the platform for more advanced use**. Leaders can make smart decisions to make advanced use easier. To give a simple example, users can get better performance by setting up versions of LLMs that are briefed on context (we explain these 'inference-time adaptation' tactics in more detail later). Leaders could take the initiative and do this instead. Background information about the company and its goals could be preloaded in a chatbot and implemented across the organisation. Users would then be given a head start for more advanced use, rather than having to create such a briefing themselves.

Capability

Operational readiness

Even when employees are motivated, organisations often fail to create the conditions that build the confidence and competence required for deeper Al use. Boston Consulting Group <u>estimates</u> that c.70% of adoption challenges stem from people and process issues rather than technical issues. Low capability emerges when organisations don't effectively reduce ambiguity, provide time for experimentation or provide opportunities to learn. This can lead to:

- → A lack of confidence: while people might feel comfortable using AI for simple tasks, they do not believe in their ability to use and apply AI in a deeper way.
- → **Cognitive overload**: people do not have sufficient mental bandwidth to focus their attention on how to use AI deeply.
- → **Ambiguity aversion**: people prefer known risks over unknown ones. Al systems can seem opaque, meaning that people might prefer not to use them, particularly for more complex tasks where there may be more risk attached. As such, avoidance (less experimentation, less feedback, less discovery) means fewer opportunities to build capability.





Leadership skills appear as a strong predictor of organisational Al adoption. Thus, ensuring the capability of leaders, alongside workers, can help to successfully integrate deeper Al adoption across an organisation.

Building operational readiness

Investment in training is likely to be part of the solution. Accenture found that 94% of workers say they want to develop new Al skills, but only 5% of organisations are providing organisation-wide training. However, here we focus on behavioural levers that can be employed alongside greater provision of training.

It is worth noting that building operational readiness requires not just providing tools and training, but also fostering the motivational environment where employees feel empowered and encouraged to develop their own capabilities.

№ What can leaders do?

- 1. **Signal institutional support**. A survey of 400 teachers **found** that exposure to credible information and strong institutional support can increase perceived usefulness and self-efficacy, thereby raising intention to adopt Al. Support needs to be specific, tailored to workers' contexts. Generic training on Al is **not useful**. Offering training at timely moments, such as when a project is about to start, might be particularly effective in encouraging adoption. Al companies have started to offer free courses (eq. Anthropic's Al Fluency course) that are designed to build capability without becoming quickly outdated.
- 2. Encourage bottom-up adoption rather than top-down. Encouraging employee-led experimentation is likely to <u>yield</u> greater gains than top-down mandates to use Al. Many firms see Al pilot projects stall because employees feel Al is being 'dumped' on them. Successful firms instead give workers agency.
- 3. **Structure the adoption journey.** Instead of expecting an immediate leap to advanced use, leaders should use 'scaffolding': a method of building competence and confidence through a series of managed, step-by-step challenges. That might start



with tasks that are low-risk and offer immediate, visible value, like encouraging employees to use AI for simple activities like summarising long documents. Then would come more complex, supervised tasks (like using AI to create a first draft of a project plan). Finally, encourage and equip employees to use AI for high-value, strategic work, like in-depth competition analysis. A UK government Copilot trial with 20,000 officials illustrates the approach: with licences, training and support, adoption stayed at ~83%, users saved ~26 minutes per day, and 82% wanted to continue beyond the pilot. Benefits spanned grades and professions, with notable accessibility gains.

Trust

Legitimacy

Lack of trust in AI tools is a significant factor inhibiting deeper AI adoption. In fact, in a <u>survey</u> of 1500+ workers in the US, lack of trust was the greatest cause for concern (45%), scoring 22pp higher than fear of job replacement (23%). A <u>further study</u> with 607 participants found a significant correlation between trust and intention to use ChatGPT.

Trust is not the same as trustworthiness. Trust is a subjective assessment, whereas trustworthiness is an objective measure of performance which can be evidenced through rigorous evaluation. Thus, leaders need to ensure that there has been an assessment of whether Al can meet their organisation's needs reliably. Only then should they deploy Al tools and address the challenge of building employee trust.

Broadly speaking, trust in technology has <u>three key drivers</u>. Presented in order of importance, these are:

- 1. **System-based factors** (expertise, predictability, and transparency);
- 2. External or environmental factors (culture, risk, and brand perception);
- 3. **User-related factors** (competency, personality traits, and workload).

The way that people assign trust to AI is not straightforward. There's evidence for both algorithmic aversion and appreciation. While <u>research</u> shows that people often penalise algorithms more harshly than humans after mistakes, <u>other studies</u> find people can develop strong preferences for algorithmic advice when tasks are objective or where algorithms have proven track records. So trust in AI is highly dependent on context and past performance.



Adding to this complexity, <u>research</u> highlights fundamental differences between how humans and AI handle information related to emotions. While humans experience emotions as deeply embodied and linked to our sense of self, AI can only analyse patterns in emotional data. Recognising these differences is critical for building appropriate trust.

▲ Improving legitimacy

What can individuals do?

1. **Increase exposure**. In our <u>ChatGOV</u> experiment, mentioned earlier, we found that all treatments that included an AI chatbot increased trust in AI by 7-13pp. Mentioning the risks of AI in the transparent bot design did not affect trust in AI. Thus, exposure to AI can itself be a method to increase trust. That could look like interactive lunch & learn sessions for tasks where AI performs consistently well.

№ What can leaders do?

- 2. Avoid Al exceptionalism in framing. While this report argues that the potential long-term impact of Al on society is exceptional and demands careful oversight (see Adapt), a key barrier to trust is the perception that Al is fundamentally different to other technologies, due to being opaque and uncontrollable. This so-called 'exceptionalism' might be the problem. As Arvind Narayanan and Sayash Kapoor argue, framing Al as a 'normal' technology allows trust to be built through familiar mechanisms such as performance, reliability and consistency over time. It can help reduce unwarranted hype on one end of the scale, and unbridled fear on the other.
- 3. **Anthropomorphise AI**. Al tools can be designed to mimic human-like cues. **An experiment** used a fictitious retail brand chatbot to test this approach. 288 Australian adults were randomly assigned to one of four treatment conditions. They were instructed to visit the fictitious website and interact with a chatbot about information for purchasing wine. Respondents then evaluated their perceptions of interacting with the chatbot, including the extent to which they thought it was human-like. The researchers found that chatbot anthropomorphism was positively correlated with



attitude towards the brand. Chatbots that were more human-like were more likely to make people think that they were interacting with another social entity. Having said this, as we discuss in Adapt, anthropomorphism can backfire and designers of Al tools should therefore be wary of making the chatbot too uncanny (ie, so human-like that its small, unavoidable flaws become unsettling to users).

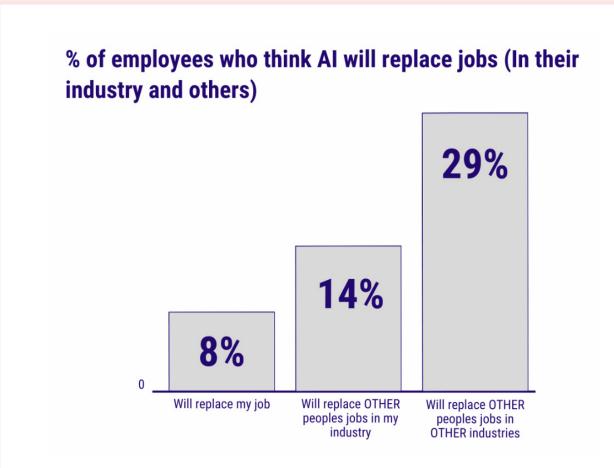
- 4. **Embed transparency**. An <u>early study (2019)</u> on Al-infused decision-making processes found that people are much more likely to trust transparent Al models than those that are 'black boxes'. Thus, providing information about how Al models were trained and tested can go some way to fostering trust. This principle aligns with the explainable Al (XAI) movement, which seeks to develop models that not only perform well but can be interpreted by humans. In a more <u>recent study</u>, researchers ran an RCT with 140 adults performing caregiving tasks in an online, simulated home healthcare environment. Individuals were randomly assigned to receive automated real-time feedback when performing their duties. They received an algorithmically determined rating. They found that real-time feedback increased perceived trustworthiness of the performance rating, which in turn improved individuals' trust in Al-generated performance ratings.
- 5. Evaluate impact and embrace the results (positive or negative). Robust evaluation can help foster a healthy 'sceptical trust', a reliance on Al that is both confident and critical. Knowing what works, particularly with high-quality evidence to back it up, can strengthen deeper adoption. It can show that, for example, there aren't unintended consequences (backfire effects) that raise ethical concerns. Alongside this, leaders should celebrate null results. Acknowledging what didn't work, and why it didn't work, can legitimise experimentation and reduce fear of failure. This creates psychological safety for further experimentation, as employees won't fear retribution if their idea doesn't work.

Threats to identity (psychological threat)

Deeper adoption may involve greater risks. That's an issue because **researchers** found that the higher the stakes or potential losses, the less people were likely to use Al. They were more likely to instinctively trust human judgement - their own or someone else's.

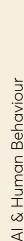


Trust in your own judgment is also bound up with concerns about how the more complex use of Al impacts your professional identity. Loss aversion means that individuals may delay adoption, fearing loss of control, loss of their livelihood, or loss of agency. Use of Al can also present a threat to your own sense of self and your self-perceived competence. The more 'personal' the task, the more that individuals are likely to be Al-averse. Interestingly, individuals also show optimism bias: they are much more likely to say that Al will pose a threat to others' jobs than their own.



Source: Irrational Labs, 2025

Self-identity is also linked to social identity, where individuals may not want to reveal they are Al users as they may be perceived to lack competence or to be lazy. In one experiment, <u>researchers</u> asked 1,215 participants to evaluate fictional employees described as receiving help from different sources. They observed a social penalty for using Al, where people who used Al were consistently rated as lazier, less competent and less diligent. This held true irrespective of the fictitious employee's gender, age or occupation.





In a follow-up experiment, they also demonstrated that those who used AI themselves were much more likely to hire candidates who used AI regularly.

In some professions, particularly those where expertise and judgement are central to professional identity (eg, healthcare), use of Al can feel like a threat. However, responses are often mixed and there are differing reactions by context, task and individual attributes. This concern also extends to the creative fields, where many artists and writers worry about Al-generated outputs. Equally, though, there are some actively using Al as a co-creative partner, allowing them to expand what's possible. Leaders should recognise this diversity in response and identify which aspects of professional identity in their organisation are most sensitive to Al.

Addressing threats to identity

№ What can leaders do?

- 1. Harness loss aversion. Loss aversion can be used to overcome reluctance to use Al. Participants in a study were asked to complete a task with either human or Al assistance. Initially, the task was framed around gains. Participants were rewarded for each correct answer. In this scenario, they showed a strong bias for human help, even when the Al was proven to be more accurate. However, the preference changed when the task was reframed around losses. In the second setup, participants started with a \$10 endowment and lost \$0.50 for every mistake. Faced with the prospect of losing money, the bias vanished. Participants in the loss scenario valued the superior Al's assistance just as much as the human's. This suggests that framing Al as a tool to prevent errors or mitigate risk (rather than just a tool for achieving gains) can make people more willing to adopt it.
- 2. **Democratise AI adoption**. Allowing employees to have a stake in AI adoption can alleviate threat concerns. That could mean using pilot programs to trial small-scale AI adoption with user feedback loops. Those who participate in the trials are likely to feel a greater sense of control. Non-participants also feel the benefits, as their peers may champion AI, reducing their own scepticism. This could even go as far as having employees **rewrite their roles with an 'AI-first' lens**. That move enables them to see AI as a positive influence on their identity, rather than as a threat.



3. **Use social proof**. Alongside democratisation, making Al usage visible and celebrated can help reduce social identity threats. Those who know someone who has used Al are three times more likely to have used Al themselves (albeit this is correlation rather than causation). This so-called 'bandwagon effect' can drive adoption at speed and scale, with minimal effort on an organisation's behalf.

Permission bias

Richard Thaler <u>coined the term</u> 'permission bias' to describe how practitioners only apply behavioural science in areas where they feel explicitly allowed to do so (ie, they've been given permission). The same principle applies to Al. Employees might engage with Al in shallow ways as they believe that is the only level of use that their organisation supports. Deeper adoption is then left untouched and even those employees who are willing and capable may not move up the adoption ladder.

▲ Removing permission bias

№ What can leaders do?

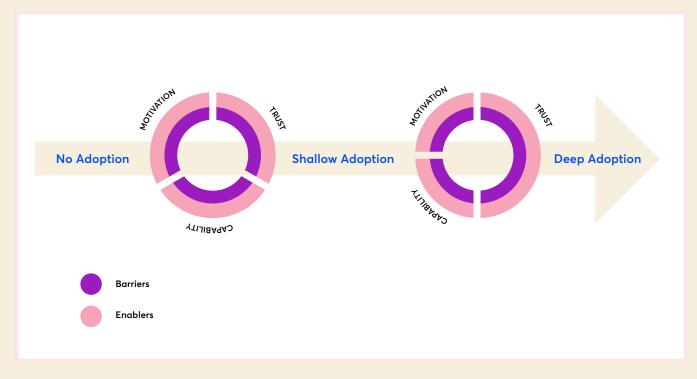
- 1. Signal clearly. Harnessing messenger effects, leaders can send unambiguous signals that AI experimentation is permissible and actively encouraged. This could be through communication, or perhaps more effectively through demoing their own use of tools. To maximise effectiveness, signals should be backed by action. That action could involve leaders modelling AI use themselves, allocating time or resources for employees to experiment, or removing obstacles that hinder exploration. Google pioneered the idea of 20% time, where workers were allowed to spend one day per week on passion projects or learning new skills. Introducing an AI-equivalent could help unlock deeper adoption.
- 2. **Use sandboxes**. Companies have begun trialling sandboxes, providing specific spaces for employees to test new Al tools. Thomson Reuters launched Open Arena in 2023, which provides a



secure internal sandbox allowing all employees to gain hands-on experience in the workplace. The creation of such an environment sends a strong signal to employees that they have permission to use Al tools.

▲ Conclusion: A roadmap to AI adoption

Below, we set out a roadmap for how organisations can move from no adoption to deep adoption:



Organisations will be at different stages on this journey for different tasks and within different teams. Tailored approaches are therefore needed. Here are the steps organisations should take:

- 1. Identify strategic priorities and high-value opportunities where Al can deliver the most impact.
- 2. Assess current adoption across teams and services and map to no adoption, shallow adoption or deep adoption.
- 3. Diagnose the specific barriers to adoption for the team/service, using the behavioural barriers outlined in this paper.
- 4. Co-design interventions with teams to encourage movement along the adoption continuum.



- 5. Pilot promising interventions, assessing take-up and effectiveness.
- 6. Scale successful interventions, taking a 'Test & Learn' approach to continuously improve and adapt.
- 7. Monitor adoption rates over time and take corrective action where uptake is stalled.

▲ How can organisations assess current adoption?

Understanding where and how AI is currently used is the first step towards improving adoption. BIT can support organisations to assess their current position using a combination of:

- → Behaviourally-informed surveys to unpack capability, motivation and trust
- → Team-level heatmaps to visualise where adoption is strongest and where support is needed
- → Usage and process data to identify adoption patterns
- → Workshops and interviews to diagnose behavioural barriers

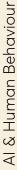
Context matters

This paper has talked about Al adoption for professional workers in general terms. Yet obviously, adoption challenges vary by:

- → Who adopts
- → Sectors
- → Types of tasks

Who adopts

There are striking disparities in who adopts. Analysing 17 studies on generative AI use, with c.140k individuals, <u>Harvard researchers</u> find that women use generative AI between 10-40% less than men. Gaps also appear across age, sector, and organisational size. Many workers report wanting to build AI skills, but receive little or no support from their employers. These figures suggest that willingness alone does not explain who uses AI and who does not.





The role and seniority of employees will, of course, also impact adoption. Psychological threat is much more likely to be present in those roles where their status, identity and influence are intrinsically linked to their expertise. Early adopters may have less to lose in terms of their identity from embracing Al. They are also more likely to be technologically literate.

Sector

Some professional sectors are more likely to see AI as challenging to their foundations than others. For example, clinicians or legal professionals may be more resistant because a significant part of their roles is based on making value judgements.

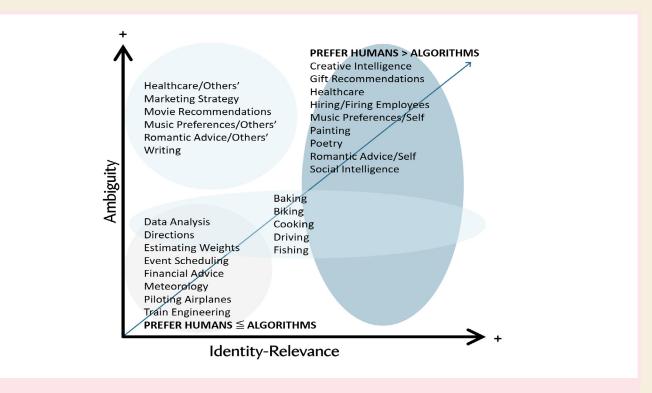
Tasks

Tasks that have high stakes or are linked to a profession's identity - for example, diagnosing a patient or approving a financial product - may require more significant efforts to encourage adoption.

One useful way of thinking about how appetite for Al varies by tasks is to map two factors onto each other:

- how central a task is to someone's self-conceptualisation (identityrelevance);
- how difficult it is to judge success (ambiguity of evaluative criteria).

The figure on the next page shows how these factors influence our stance on AI across some popular tasks:



Source: Decisions with algorithms, 2025

Design matters

Adoption can also hinge on factors outside an organisation's direct control. In particular, the design of the tools themselves can heavily influence adoption. Wharton's 'Blueprint for effective Al chatbots' synthesises a range of studies showing how the design of chatbots influences user experience and trust. Findings include:

- → Labelling AI as 'learning': users were more likely to trust and rely on chatbots when they were framed as improving over time.
- → Framing the AI as superhuman: non-technical users responded better to tools positioned as uniquely capable, rather than as replacements for human agents.
- → Allowing modification and control: users were more satisfied and willing to engage when they could customise their interactions with a chatbot.
- → Avoiding overly human-like features in high-pressure contexts: for example, machine-like bots were better received when delivering bad news or collecting sensitive personal data.

While these studies are focused on customer satisfaction and sales, lessons from these findings can equally apply to adoption among employees. Organisations should choose AI products whose features match the needs of the context.