
The Psychology of the AI Age

Human Behaviour, Identity,
and Decision-Making

A Research Synthesis for the Era of Generative Systems

Contents

Preface	xv
1 The New Psychological Frontier	1
1.1 A Species That Thinks With Its Tools	1
1.2 From Calculator to Interlocutor: What Actually Changed	3
1.3 A Short History of a Long Relationship	4
1.4 Four Layers of Psychological Impact	6
1.5 Co-Evolution, Not Replacement	10
1.6 What the Evidence Can and Cannot Yet Tell Us	11
1.7 How This Book Is Organised	12
2 Cognitive Offloading and Executive Function	15
2.1 The Economics of Mental Effort	15
2.2 The Offloading Instinct: From Notebooks to Neural Networks	16
2.3 The Cognitive Offloading Cycle	18
2.4 Memory: What We Keep When the Machine Remembers	19
2.5 Attention and the Fragmented Workspace	21
2.6 Critical Thinking and the Metacognitive Illusion	22
2.7 Skill Acquisition and the Deskillng Question	24
2.8 What Protects Cognition	26
3 Identity, Agency, and the Synthetic Self	29
3.1 The Self as an Ongoing Construction	29
3.2 The Extended and Digital Self	31
3.3 Authorship, Ownership, and the Ghostwriter Effect	32
3.4 Agency: The Feeling of Being the Cause	35
3.5 The Synthetic Identity Model	37
3.6 Algorithmic Self-Knowledge and the Shaping of Preference	39
3.7 Two Trajectories: Integration and Diffusion	40

4	Decision-Making in the Algorithmic Loop	43
4.1	Bounded Rationality Meets Unbounded Advice	43
4.2	Trust, Reliance, and Their Miscalibration	45
4.3	Automation Bias and Complacency	46
4.4	Aversion and Appreciation: A Contingent Story . . .	47
4.5	The Human–AI Decision Matrix	48
4.6	Choice Architecture, Hypernudging, and Manipulation	51
4.7	Moral Judgement in the Loop	53
4.8	Designing for Appropriate Reliance	54
5	Social Mechanics, Empathy, and Attachment	57
5.1	Mindless Sociality: Why Machines Get Treated as People	57
5.2	Anthropomorphism: Mechanism and Individual Dif- ference	59
5.3	The Uncanny Valley and the Limits of Likeness . . .	60
5.4	Disclosure Without Judgement	61
5.5	The Attachment Gradient	63
5.6	Loneliness: Substitution or Scaffold?	65
5.7	AI-Mediated Communication and Human-to-Human Relations	67
5.8	Vulnerability, Dependence, and the Duty of Design .	68
6	Well-Being, Anxiety, and Existential Dread	71
6.1	Naming the Anxiety	71
6.2	Measuring What People Feel About AI	73
6.3	Work, Identity, and the Threat of Displacement . . .	74
6.4	Technostress and the Demands–Resources Account .	77
6.5	Pathways to Well-Being	78
6.6	Meaning, Mortality, and the Existential Register . . .	80
6.7	AI in Mental Health Care: Promise and Guardrails . .	81
7	Societal and Organisational Transformation	85
7.1	From Individual Cognition to Institutional Cognition	85
7.2	Algorithmic Management and the Contested Workplace	87
7.3	The Productivity Evidence and the Jagged Frontier . .	89
7.4	Task, Role, Organisation, Institution	91
7.5	Education: Assessment, Integrity, and Learning . . .	93
7.6	Governance, Explanation, and Accountability	95

7.7	Fairness, Bias, and Institutional Trust	96
8	Synthesis and the Human Future	99
8.1	What We Now Know	99
8.2	The Framework Revisited	102
8.3	Cultivating Cognitive Resilience	105
8.4	Designing for Agency	108
8.5	Governing for Human Flourishing	109
8.6	The Cognitive Stewardship Framework	111
8.7	A Research Agenda for the Next Decade	112
8.8	Coda: The Human Future	114
A	Measurement Instruments for Human–AI Psychological Research	117
A.1	Instruments in Current Use	117
A.2	Constraints Worth Stating Explicitly	117
B	Protocols for Practitioners	121
B.1	For Educators	121
B.2	For Managers Deploying AI	122
B.3	For Clinicians	123
	Glossary	125
	References	129
	Index	157

CHAPTER 1

The New Psychological Frontier

A team leader drafting a difficult message, a clinician phrasing a diagnosis, a student framing an argument: in each of these moments a conversational system can now supply the words, and increasingly it does. This chapter sets out why that matters psychologically. It places the phenomenon inside a long tradition holding that human thinking has always been distributed across bodies, artefacts and other people, so that the useful question is never whether people offload but what changes when the receiving artefact is general, conversational and generative. It traces the arc from the imitation game through ELIZA to the systems released to the public in 2022, sets out the four-layer framework and the three moderators that organise the book, and states candidly how much of the present evidence is cross-sectional, self-reported and short-horizon rather than settled.

1.1 A Species That Thinks With Its Tools

On a Tuesday evening a team leader in a mid-sized logistics firm has to tell a colleague of nine years that her contract will not be renewed. She writes two sentences, deletes them, and then opens a chat window and asks for help with the tone. Ninety seconds later she has a message that is warmer than hers, better organised than hers, and which she sends after changing three words. Nothing in that sequence is dramatic. She did not abdicate a decision; the decision had been made elsewhere. What she delegated was the part of the task that carried the moral weight: the choice of how one human being tells another that something has ended.

Scenes of this kind are now unremarkable, and their ordinariness is precisely what makes them worth theorising. The temptation is to treat them as a break in human history, but the more defensible starting point is continuity. Cognitive science spent the closing decades of the twentieth century dismantling the idea that

thinking happens inside a single skull. Hutchins (1995) showed, by watching navigation teams bring a warship into harbour, that the computational work of fixing a position was carried out by an assembly of people, charts, alidades and standardised procedures, no part of which held the whole task. Clark and Chalmers (1998) pressed the argument philosophically: when a notebook plays the functional role that biological memory would otherwise play, and does so reliably and transparently, there is no principled reason to place the boundary of the cognitive system at the skin. Clark (2003) went further, describing humans as natural-born cyborgs whose distinguishing talent is not raw processing power but a willingness to annex whatever scaffolding is available.

This tradition also supplies a vocabulary for the mechanics of offloading. Kirsh and Maglio (1994) distinguished pragmatic actions, which move a task towards its goal, from epistemic actions, which rearrange the world so that less thinking is required — rotating a falling Tetris piece rather than mentally imagining the rotation. Rearranging the world to reduce mental work is not a modern vice; it is the signature manoeuvre of a species with a small working memory and a large capacity for building things. Hollan et al. (2000) turned this into a research programme for human–computer interaction, arguing that interfaces should be studied as parts of larger cognitive systems rather than as channels to an isolated mind. Most usefully here, Salomon et al. (1991) drew the distinction between the effects *with* a technology, meaning what a person can do while working in partnership with it, and the effects *of* a technology, meaning the cognitive residue left in the unaided mind afterwards. Two people can show identical performance with a tool and diverge sharply in what remains when it is taken away, and the second quantity is the one this book is chiefly about.

The consequence for this book is a discipline of question-asking. Asking whether cognitive offloading is good or bad is a category error, roughly like asking whether writing is good or bad; Risko and Gilbert (2016) note that offloading is better understood as an adaptive strategy whose value depends on the fit between the demand, the tool and the person's ability to judge that fit. The question worth asking is what is different now, and answering it requires precision

about the new artefact rather than gestures at its novelty.

1.2 From Calculator to Interlocutor: What Actually Changed

Four differences seem to us to carry most of the psychological weight, and none of them concerns raw capability. The first is generality. A calculator, a spreadsheet, a spelling checker and a search engine are domain-bound: each has a mouth-shaped opening into which only certain problems fit, and the user must formulate the problem before the tool can help. Systems built on large pretrained models are, by design, not bound in this way; the argument of Bommasani et al. (2021) in naming them foundation models was precisely that a single trained artefact could be adapted to an open-ended range of downstream tasks. Whether a given system is good at a given task is a separate and much messier matter, but the user no longer has to decide in advance which tool a problem belongs to. Everything can be brought to the same window.

The second is that the interface is conversational rather than command-driven. A command language requires the user to hold a model of what the system will accept; a conversation requires almost nothing, which is part of why adoption ran ahead of any organised programme of training. Formulating the request in ordinary words changes what gets delegated, because vague, value-laden and half-formed problems can now be handed over in the state in which people actually have them.

The third difference is that the systems are generative rather than retrieval-based. A search engine returns documents that existed before the query and leaves synthesis to the reader; a language model returns a fluent artefact that has never existed before, indistinguishable in form from something a competent colleague would have written. As Shanahan (2024) argues, it is important not to slide from that fluency into attributing beliefs and intentions to a system whose operation is sequence prediction; but the psychological point stands regardless of how we describe the mechanism. The output arrives finished, and finished text invites assent rather than evaluation.

The fourth difference is the one that organises the rest of this book. Earlier cognitive technologies took over storage, transmission and

calculation. These systems participate in judgement and expression: what to conclude, how to phrase it, which of two options is more reasonable, what to say to a grieving friend. The technical reports accompanying the frontier models of 2023 documented performance across professional examinations and open-ended reasoning tasks that few observers had expected so soon (OpenAI et al., 2023), and the early assessment by Bubeck et al. (2023) went as far as describing sparks of general intelligence, a claim that should be read with the caution appropriate to an unrefereed report written by the model developer's collaborators using non-public evaluation conditions.

Yet one does not need to settle the question of machine intelligence to observe the human consequence. Epstein et al. (2023) note that generative systems intervene in creative practice at the level of authorship rather than production, unsettling the attribution of credit that creative fields run on. And the participation is measurable in the scholarly record itself: Liang et al. (2025) estimate, using a distributional method applied to hundreds of thousands of papers, that a substantial and rising fraction of sentences in recent abstracts and introductions bear the statistical marks of model-assisted writing, with computer science well ahead of other fields. Scientists writing about the tool were among the first to write with it.

1.3 A Short History of a Long Relationship

None of this arrived without warning. Turing (1950) set aside the question of whether machines can think as too ill-defined to answer and replaced it with a behavioural test built on conversation, thereby fixing the terms of the argument for seventy years: dialogue would be the arena in which machine intelligence was contested. A decade later Licklider (1960) sketched a different aspiration — not a machine that imitates a person but a coupling in which the human formulates goals and hypotheses while the machine handles the routine work of preparing the way for insight. His essay is an early and unusually clear statement of the partnership framing that still dominates practitioner discourse, and its optimism rested on an assumption this book treats as open: that the division of labour between partners would remain stable.

The first evidence that it would not come almost immediately, and by accident. Weizenbaum's ELIZA was a pattern-matching program of a few hundred lines whose DOCTOR script reflected users' statements back as questions; Weizenbaum (1966) reported that users conversed with it at length and attributed understanding to it, some of them resisting his explanations of how few lines of code were involved. What later became known as the ELIZA effect is an early documented case of people responding as if to a mind while knowing, propositionally, that there was none.

Three decades later Nass et al. (1994) turned this from an anecdote into an experimental literature, showing that people apply politeness norms and reciprocity to computers, and treat cues such as voice as social signals, without believing for a moment that computers are people; Reeves and Nass (1996) generalised the result as the media equation: social responses to media are automatic, and knowing better does not switch them off.

Meanwhile a parallel literature was accumulating on the cognitive side of the same relationship. Tversky and Kahneman (1974) had established that human judgement runs on heuristics with systematic biases, which is the precondition for any argument about when advice helps. Bainbridge (1983) observed that automating the easy parts of a task can leave the human operator less able to handle the hard parts, a finding whose relevance to knowledge work is only now being taken seriously. Sparrow et al. (2011) showed that participants who expected continued access to information remembered the information less well while remembering better where to find it, and Risko and Gilbert (2016) eventually gave the whole phenomenon a name and a framework. The 2022 release of conversational language models to the general public therefore did not open a new field so much as arrive in the middle of several existing ones, whose findings were suddenly load-bearing. Figure 1.1 sets out this arc, and the lag it displays is itself part of the argument: the psychological questions have usually been posed after the capability arrived rather than before it.

One shift in stance is needed before the empirical chapters can proceed. Rahwan et al. (2019) argue that intelligent systems are now too complex, adaptive and widely deployed to be understood

Milestones in the Human-Machine Cognitive Relationship




Figure 1.1. Milestones in the study of the human-machine cognitive relationship. The three colour-coded tracks show that findings about human cognition and human emotion have consistently trailed, rather than anticipated, jumps in machine capability.

by inspecting their code, and must be studied with the methods of behavioural science: observation in the wild, experiment, and analysis of the hybrid human-machine ecologies they inhabit. This book takes the symmetrical position. What such systems do to human psychology also cannot be read off their architecture; it has to be measured where people actually use them. Measurement of that kind needs a map, because one act of delegation registers at several levels at once — in a mind, a self-concept, a relationship and an institution — and it is to that map that the chapter now turns.

1.4 Four Layers of Psychological Impact

The literature on human responses to artificial intelligence (AI) is large, fast-growing and badly organised. Findings about memory sit next to findings about loneliness, and studies of authorship next to studies of hiring algorithms, with little sense of how they bear on one another. This book imposes a structure of four layers, arranged from the interior of the individual outwards, and that structure is a claim rather than a convenience.

The first is the cognitive layer: attention, working memory,

executive function, and the decision to offload. This is where the mechanics of delegation live, and where the classical literature on external memory and cognitive offloading gives us the sharpest tools (Risko and Gilbert, 2016; Sparrow et al., 2011). The second is the identity layer. Delegating the composition of a sentence raises a question that no calculator ever raised, namely whose sentence it is; and experimental work on the sense of agency — the ordinary background feeling of being the author of one's own actions — indicates that this feeling is inferred and manipulable rather than a direct readout of one's own causal contribution (Haggard and Chambon, 2012). That is a narrower matter than agency in the sense Bandura (2006) gives it, the capacity to set goals and act on them, but the two meet wherever fluent assistance is available: studies of co-writing already show that users can fail to feel ownership of text they nonetheless sign (Draxler et al., 2024).

The third is the social-affective layer, which covers what happens when the system is treated as a social partner: anthropomorphism, trust, empathy, attachment. Social responses to machines are mindless rather than credulous, triggered by cues rather than by beliefs about what the machine is (Nass and Moon, 2000), and some users form relationships they describe in the vocabulary of friendship (Brandtzaeg et al., 2022). The fourth is the societal layer: the organisations, classrooms and governance regimes in which individual delegation is aggregated, priced, mandated or forbidden, and in which algorithmic systems become instruments of control rather than merely of assistance (Kellogg et al., 2020).

Judgement is treated here as a cross-cutting process rather than as a fifth layer. Deciding whether to accept an answer draws on the executive resources that govern offloading at the cognitive layer and on the social heuristics that a fluent interlocutor recruits at the social-affective layer, so trust and reliance run across both of those layers rather than sitting beside them. Chapter 4 is where that process is examined.

Two further elements complete the framework, shown in Figure 1.2. The first is a set of three moderators, and they carry the book's central argument. Task structure matters because effects invert with the fit between the task and the system's competence: in a field experiment

with several hundred consultants, access to a frontier model raised quality on tasks inside its capability while degrading it on a task deliberately placed outside (Dell'Acqua et al., 2023), a result reported in a working paper that has not been through peer review.

Individual differences matter because a tool that supplies finished answers meets an expert and a novice differently: expertise is accumulated through effortful, feedback-rich practice (Ericsson et al., 1993), so the person who already holds an internal model can audit the output, while the person still acquiring one may be bypassing the practice that would let them audit it later. Institutional design matters because whether a delegation is followed by verification depends more on how the work and the interface around it are arranged (Amershi et al., 2019) than on the user's intentions, and because who is answerable for the result is a question about the organisation rather than the individual (Santoni de Sio and van den Hoven, 2018).

The second element of the framework is the fork the moderators produce: augmentation, in which capability expands while agency and competence are preserved, or erosion, in which skills atrophy and responsibility diffuses until no one in the loop is quite the author of the outcome. Neither branch is the natural destination of the technology, and the difference between them is the subject of this book.

The layers are analytically separable but causally entangled. Return to the team leader and the redundancy message. At the cognitive layer she has offloaded a composition task, sparing herself the effort of finding words and, if the offloading literature reviewed in Chapter 2 generalises to composition, some of the encoding that comes from having found them. At the identity layer she now occupies an unstable position: she is the sender of a message she did not write, and if asked whether it expressed her own regret she would have to think before answering.

At the social-affective layer the exchange was easy precisely because the system is a listener that does not judge, and asking it for help with an unpleasant task carries none of the interpersonal cost of asking a colleague. At the societal layer, if her firm comes to expect communication of this quality and speed, the delegation ceases to be a choice; and the colleague receiving the message is now a person




Figure 1.2. The four-layer framework of human–AI psychological interaction that organises this book. Reading outward from the human and the system at the centre, each layer names a domain of documented psychological effect; the moderators along the base determine which of the two divergent outcomes on the right a given person, team, or institution reaches. Judgement, trust and reliance are treated as a process cutting across the cognitive and social-affective layers rather than as a fifth layer, and Chapter 4 examines it.

whose dismissal was phrased by a machine. One act, four registers. The layers are separated in this book to make the evidence tractable, and reassembled in Chapter 8 because that is how they operate.

1.5 Co-Evolution, Not Replacement

The dominant popular framing is substitution: the machine takes over a task and the human is left with less to do. It is the wrong frame for psychology, because it describes a transfer of work rather than a change of relationship. The more accurate frame is co-evolution. People adapt to the systems, the systems are trained on what people produce, institutions adjust their expectations to what has become easy, and the resulting equilibrium is a property of the whole arrangement rather than of either party. Sundar (2020) makes the theoretical case for treating human–AI interaction as its own object of study, on the ground that machine agency introduces psychological dynamics — cues, heuristics, attributions of authorship — that neither media psychology nor human factors alone anticipates. Guzman and Lewis (2019) argue in parallel that communication scholarship must expand from studying communication *through* machines to studying communication *with* them.

This is also why the field's most successful predictive models are of limited use here. The technology acceptance tradition, from Davis (1989) through the unified account of Venkatesh et al. (2003), explains uptake with impressive economy: perceived usefulness, perceived ease of use, social influence, facilitating conditions. They answer who adopts a system and how quickly, but were never designed to answer what adoption does to the adopter; and where use has already become routine, variance in intention to use is a poor place to look for psychological consequence. The interesting variance has moved downstream, into what people do after adopting and what remains of their capability afterwards.

Co-evolution also implies a design choice rather than a trajectory. Brynjolfsson (2022) names the trap precisely: because human-like performance is the most legible benchmark, effort and capital flow towards systems that substitute for people rather than towards systems that extend what people can do, even though the second path generally produces more value and distributes it more widely. The

trap is not technological determinism but the opposite — a decision, made repeatedly and often without deliberation, about what to build and what to reward.

1.6 What the Evidence Can and Cannot Yet Tell Us

A book with this subject matter must be candid about the state of its evidence, because the confidence of public discussion far exceeds the strength of what has been established. Four limitations recur, and readers should hold them in mind whenever a finding in later chapters sounds decisive.

The first is design. A large share of the available work on generative systems and thinking is cross-sectional and self-reported. Gerlich (2025) surveyed several hundred participants and found that frequent reliance on such tools was associated with lower critical-thinking scores, with younger respondents scoring lowest. The pattern deserves attention, but it is a cross-sectional association: the arrow could run either way, or through a third variable such as a prior disposition to effortful thought. Lee et al. (2025) asked knowledge workers about hundreds of real instances of generative tool use and found that greater confidence in the system went with less reported critical engagement, while greater confidence in one's own expertise went with more. That is among the most informative descriptive evidence available, and it still rests on workers' retrospective reports of their own mental effort, which people are not reliably able to observe in themselves.

The second is sampling. The participants in this literature are overwhelmingly university students and white-collar knowledge workers in wealthy countries, recruited through course credit or online panels, performing short tasks in unfamiliar settings. The populations for whom the stakes are highest — clinicians, teachers, benefits assessors, people whose jobs are being restructured around these systems — appear far less often, so the field's confidence about working life rests on its least representative samples.

The third is effect size and analytic freedom. The nearest precedent for the present debate is the argument over screens and adolescent wellbeing, and its lesson is chastening. When Orben and Przybylski (2019) applied specification-curve analysis to three large datasets, test-

ing the enormous number of defensible analytic choices rather than one, the association between digital technology use and adolescent wellbeing proved negative but negligible, comparable in magnitude to associations with wearing glasses or eating potatoes. Przybylski and Weinstein (2017) had already shown that the relationship was not even monotonic: moderate use was associated with slightly better wellbeing than none. Claims about generative systems are currently being made with the same confidence and less data.

The fourth is speed, in two senses. Findings now circulate as preprints long before review; the widely publicised electroencephalography study of essay writing with an assistant (Kosmyna et al., 2025) was reported in the press as evidence of cognitive damage on the basis of a small, unrefereed study whose own authors' framing was more cautious than the coverage. And the systems change faster than good longitudinal work can be completed, so that a three-year panel study necessarily measures a tool that no longer exists.

What the field has established is narrower than it looks, but not nothing. Surveying a large body of empirical studies of human–AI decision making, Lai et al. (2021) found a literature dominated by short, artificial tasks with proxy participants and inconsistent outcome measures, making cumulative conclusions difficult. And the most systematic answer to the partnership question is deflationary: pooling more than a hundred experiments, Vaccaro et al. (2024) found that human–machine combinations on average performed worse than the better of the human or the system alone, with losses concentrated in decision tasks and gains largely confined to tasks involving the creation of content. Synergy is possible, but on this evidence it is an achievement of design rather than a default, and the pooled estimate is assembled from short laboratory studies of the kind just criticised.

1.7 How This Book Is Organised

The chapters follow the four layers, moving outward from the individual mind and closing with the institutions that shape it. Chapter 2 opens the cognitive layer with offloading, memory and executive function. It sets out a cycle running from task demand to internal encoding, identifies the point at which reinvestment of freed capacity diverges from attrition, and reads the search-engine

literature as a warning about metacognition: externalising knowledge inflates our estimate of what we know unaided (Ward, 2021).

Chapter 3 takes up identity and agency. It treats the self as an ongoing construction rather than a stored object, and asks what happens to authorship, self-efficacy and narrative coherence when a fluent system participates in the composition of one's own words and opinions (Draxler et al., 2024; Jakesch et al., 2023). Chapter 4 then turns to judgement, trust and reliance, the cross-cutting process that runs across the cognitive and social-affective layers rather than forming a layer of its own. It separates trust as an attitude from reliance as a behaviour, argues that the design goal is calibration rather than confidence, and shows why algorithm aversion and algorithm appreciation are not contradictory findings but conditional ones (Dietvorst et al., 2015; Logg et al., 2019).

Chapter 5 opens the social-affective layer, asking why machines are treated as partners at all. It examines mindless sociality, the three-factor account of when and why people attribute human qualities to non-human agents (Epley et al., 2007), and the disclosure advantage that a non-judging listener enjoys (Lucas et al., 2014). Chapter 6 continues at that layer but shifts to consequence: loneliness, attachment, anxiety about displacement, technostress, and meaning. The evidence here is genuinely divided, with careful work reporting both relief from loneliness and harm from dependence (De Freitas et al., 2025; Laestadius et al., 2022).

Chapter 7 raises the level of analysis to organisations, education and governance, treating firms as cognitive systems and algorithmic management as a contest over control, and separating the different speeds at which task, role, organisation and institution adapt. Its central claim is that measured aggregate productivity gains and measured individual capability loss are observations at different horizons rather than rival results: output rose in a staged rollout of an assistant to call-centre agents (Brynjolfsson et al., 2025), while in a randomised classroom trial students who had practised with a standard chatbot did worse on a later unaided test than peers who had never had it (Bastani et al., 2024), a second working paper that has not been through peer review. Chapter 8 then reassembles the layers, states plainly which findings are robust enough to build policy on,

and sets out what individual habits and institutional arrangements would have to look like for the augmentation branch to be the likely one (Shneiderman, 2020b; Santoni de Sio and van den Hoven, 2018).

We begin where the mechanism is clearest and the evidence longest. Long before conversational systems existed, psychologists had established that people offload eagerly, that they are poor judges of what they have retained, and that the feeling of understanding can survive the loss of the understanding itself. Chapter 2 asks what those findings become when the external store no longer holds facts but produces answers.

Chapter in Brief

- Psychological effects of conversational systems are best organised into four entangled layers — cognitive, identity, social-affective and societal — rather than treated as a single undifferentiated impact.
- A single act of delegation can register on all four layers at once, so a study confined to one layer may misestimate what the delegation has actually done.
- Task structure, individual differences and institutional design are the moderators this book argues decide whether delegation yields augmentation with preserved agency or erosion through skill atrophy and diffused responsibility.
- Technology acceptance models explain who adopts these systems and how fast, but were never designed to explain what adoption does to the adopter's competence or sense of authorship.
- Much current evidence is cross-sectional, self-reported and drawn from students and knowledge workers, and pooled results show human-machine teams often underperforming the better partner alone.