HUMAN ADAPTIVE INTELLIGENCE

THE PILLARS OF MUTUALITY AND TRUST



A PROPOSED SHIFT IN THE CONCEPT OF INTELLIGENCE — WHAT PILOT TRAINING, QF32 AND 1,500 DOCUMENTED AI TRANSITIONS SAY ABOUT THINKING WELL TOGETHER

DENIS JENKINS | MATTHEW BYRNE

SEPTEMBER 2026 · © DENIS JENKINS AND MATTHEW BYRNE 2026

IN ONE SENTENCE

Human Adaptive Intelligence — the capacity of a trusted network of minds to construct, test and refine models of reality together — rests on pillars that can be named, taught and measured: language and information management, the shared disciplines of crew resource management, and an unbroken connection between the person doing the work and the decision. Across 1,500 documented AI transitions, the organisations that kept those pillars intact avoided the backlash. This paper sets out what the pillars are, how the evidence confirms them, and what a third paper must do next.

By Denis Jenkins and Matthew Byrne — Mutuality — The Future of Trust Project. The second paper in a series: it builds on *Mutuality and Trust Are the Foundation of Human Adaptive Intelligence* (Jenkins, 2026) and draws on the Fifth Revolution evidence bank at N=1,500 (*The White Whale Report 2026*), *The 5th Revolution* (Byrne, 2026) and *Mutuality: The Future of Trust* (Byrne and Jenkins, 2024). Written for the people who have to decide — a starting point, not a final answer. The five terms the paper turns on are defined on page 5.

Eight insights for practitioners

1 **Mutuality has a cognitive floor.** People cannot reliably read one another's intentions, values or competence without the ability to acquire, organise, interpret, communicate and evaluate information. Language and information management are not communication skills; they are the infrastructure on which trust is built — and, in Jenkins's three decades of profiling, they can be developed at any age.

2 **The chain runs both ways.** Weak information management leads to misinterpretation, misunderstanding, uncertainty, reduced trust, weaker mutuality and diminished adaptive intelligence. Greater trust leads to more information shared, better understanding and more trust. Which loop an organisation is in is decided before the unknown arrives.

3 **The record shows the chain at scale.** Across 1,500 AI transitions, every half-point that AI acceleration (the engine) runs ahead of human capability (the driver) raises the share of cases ending in backlash: 3% at half a point, 11% at one, 53% at one and a half, 81% at two (r = 0.82 with consequence intensity). The gap, not the tool, is the risk.

4 **Same technology, different design.** In the 305 cases (20.3%) where people were enabled ahead of the machine, none ended in backlash and none reached a severe consequence. In the other 1,195, one in three ended in backlash and one in five reached a severe consequence.

5 **Crew Resource Management is the green line's playbook.** The three moves that recur in the successful cases — change negotiated rather than announced, governance in place before scale, capability built at rollout — are CRM's disciplines translated out of the cockpit, and cases carrying them show uplift at about twice the rate of the rest. The order is Mutuality, Culture, Capability, then Integration.

6 **Connection is accountability.** Human Adaptive Intelligence flourishes when the person performing a function stays cognitively connected to observing, evaluating and improving it. In the record, accountability was located by design in 84% of enterprise uplift cases and 91% of the 45 that held through open conflict, and in none of 80 backlash cases. A person who cannot change a decision cannot be accountable for it.

7 **Mattering decides whether cognition is offered.** Beneath the pillars runs an older question — *do I still matter here?* Legitimacy, the bank's nearest measure of trust, is the strongest early warning it holds: among design-window cases scoring 3.5 or higher, 35% show uplift; below that, 3%.

8 **Augmentation, not automation, is the test for AI.** Assistance built around a person doing a job lifts the least experienced most; capability dropped on people as advice widens the gap. AI that helps a worker observe, compare, question, model, test, refine, understand and improve returns process intelligence to the person doing the work. That test is the agenda of the third paper.

CONTENTS

Inside this paper

In one sentence · Eight insights for practitioners 2

Reading key: the five terms this paper turns on 4

1 · The cockpit: when the mental model failed, the crew became the intelligence 5

2 · The cognitive pillars: mutuality has a cognitive floor, and it can be raised 6

3 · The chain: from information management to trust — and the whale at scale 8

4 · Shared cognition: Crew Resource Management, the green line and the order of operations 11

5 · Connection and accountability: who can change the decision? 13

6 · Emotional energy, mutuality and mattering: the trust bank has a clock on it 15

7 · Automation or augmentation: what the evidence studies say 18

8 · The four capabilities and the cognitive foundations 20

9 · The turn to human intelligence 21

10 · Towards the third paper: five questions 22

The invitation: together is better 24

About the research and the authors 25

References 27

How to use this paper. Sections 1 and 2 set out the cognitive pillars observed in pilot training and on QF32. Sections 3 to 8 test them against the Fifth Revolution evidence bank and the two books. Sections 9 and 10 record the wider turn to human intelligence and set the agenda for the third paper. Every figure from the evidence bank is cited by exhibit in *The White Whale Report 2026* and resolves to Frozen Package N1500 v1; every case named carries its identifier.

Reading key: the five terms this paper turns on

Two vocabularies meet in this paper. Jenkins writes of Human Adaptive Intelligence, HAI, the habitat: the trusted network in which intelligence lives. Byrne writes of Human Intelligence, HI, as a transition capability: what each person brings into that habitat and grows within it. The terms below are used in both senses without apology; where the evidence bank scores "HI transition capacity", it is scoring something Jenkins would recognise at once.

Human Adaptive Intelligence

The capacity of individuals and communities to continually improve their understanding of reality and their ability to act effectively through the interaction of cognitive development, mutuality and trust, reflection, structured thinking, collective learning and, increasingly, artificial intelligence (Jenkins, 2026). Intelligence is no longer simply a property of a brain; it becomes a property of an ecosystem.

Information management

The ability to acquire, organise, interpret, communicate and evaluate information effectively. In Jenkins's account it is the cognitive infrastructure beneath language, mutuality and trust: without it a person cannot reliably determine another's intentions, values or behaviour. Developed through recognition, categorisation, sequencing, comparison and pattern detection.

Mutuality

A tool, a process and a practice that builds respect and trust (Byrne and Jenkins, 2024); operationally, a power-sharing practice in which the people who will use, be affected by or carry the consequences of a system are treated as co-authors rather than users, realised through voice, explainability and override (Byrne, 2026). Not consultation theatre.

Trust and legitimacy

Trust is the residue of repeated experiences of being noticed, valued and relied upon. The evidence bank cannot measure it directly; its nearest measure is legitimacy: whether the people affected found the way AI was introduced fair and worth their trust, scored 0–5 from the public record.

The engine and the driver

AI acceleration is the engine: how fast and how widely a capability is deployed. HI transition capacity is the driver: the judgement, skills, contestability, role clarity and decision authority of the people expected to work with it. Uplift is a case in which the driver scored above the engine: 305 of 1,500 cases, 20.3%.

How the evidence is read — and what it is not. The Fifth Revolution evidence bank is a purposively assembled, source-referenced collection of documented AI transitions, not a random sample. Each case carries five 0–5 scores, a phase and an outcome path, coded from the public record and independently double-coded on a blinded sample. Every contrast in this paper is a descriptive difference between coded groups, not a causal estimate; organisations that enabled their people first may differ in other ways. What the record shows is where the outcomes sit, and how consistently.

When the mental model failed, the crew became the intelligence

The case

QF32, an Airbus A380, left Singapore for Sydney on 4 November 2010 with 469 passengers and crew. Minutes after take-off an engine failed uncontained, damaging systems across the aircraft and flooding the cockpit with fault messages (ATSB, 2013; de Crespigny, 2012).

On that day there were five pilots in the cockpit rather than three. Captain Richard de Crespigny was under annual route check; alongside First Officer Matt Hicks and Second Officer Mark Johnson sat Senior Check Captain Dave Evans and Check Captain Harry Wubben. Before departure de Crespigny briefed his crew to focus on keeping the flight safe and not to stay quiet or become distracted because a check was under way. The hierarchy was in place and running smoothly. Then the engine exploded.

De Crespigny concentrated on stabilising the aircraft and directed the first officer to the ECAM fault notices. What arrived was a mass of failures too large for any protocol, and the most prepared aviator in the cockpit reached the edge of individual cognition. He later described the moment precisely: "All these checklists coming in were filling up my mental model. I'd lost all my free mental space. I couldn't absorb more failures, and I'd lost the ability to create the complex knock-on effects in my mind. My mental model of the aircraft had failed."

This was not a failure of preparation. De Crespigny had flown gliders and fighters, studied the manuals, aerodynamics and structural design of every aircraft he flew, and held that automation breeds complacency in anyone who does not understand the why behind it. That is the conventional picture of intelligence at its best: the capacity to visualise a complete system, recognise its dependencies and hold a complex model in mind. And it ran out. The information model had become too large for one brain.

What saved the aircraft was a deliberate shift. The rules had changed; the aircraft no longer behaved as the manuals described; and the crew moved from rules and protocols to principles and practices, from command hierarchy to a flatter, shared process in which responsibility for parts of the problem was distributed and the findings were drawn back into collective decision-making. The two check captains, who outranked de Crespigny, set rank aside and worked as equals toward a single goal. Not one of them could have dealt with the crisis alone. Jenkins's first paper named what they became: a collective intelligence he calls Human Adaptive Intelligence. De Crespigny's own summary is the shortest statement of it: "I knew we had the cognitive tools and the brains on board to solve this."

> *"Most theories of intelligence begin with the individual and ask how groups benefit from intelligent people. This theory begins with relationship and asks how relationships increase the intelligence of individuals and communities." — Jenkins (2026)*

That inversion is the foundation on which this paper builds. Adaptive intelligence, on Jenkins's account, is not confined to the individual mind; it is an emergent property of a trusted network of minds that can collectively construct, test and refine models of reality. It becomes necessary precisely when the information model is too large for one person, and it makes trust not a moral ideal but a practical requirement: mutuality lets responsibility for parts of a problem be distributed; trust lets those parts be recombined into a coherent response. The question this paper asks is what that network is made of — and what happens to it when the unknown that arrives is not an exploding engine but artificial intelligence.

Mutuality has a cognitive floor — and the floor can be raised

Where the observation comes from

Jenkins has profiled commercial pilot candidates for many years. Low cognitive scores fared poorly in simulator training; high scores did well; and trainees who missed the grade at first succeeded after deliberate cognitive development. Pilot training is therefore the benchmark for where mutual trust and Human Adaptive Intelligence flourish.

A recurring observation across three decades of profiling is that human mutuality requires a minimum baseline of cognitive capability. The baseline is language and information management: the ability to acquire, organise, interpret, communicate and evaluate information. Without it, a person cannot reliably determine the meaning of another person's intentions, values or behaviour, and concepts such as shared values, integrity and trust cannot be accurately established or sustained. Language and information management are therefore not merely communication skills. They are the foundational cognitive infrastructure on which mutuality, trust and Human Adaptive Intelligence are built.

Pilot training is chosen as the benchmark because it is an unusually systematic example of cognitive development. A competent commercial pilot is not simply taught to fly an aircraft; they are trained to develop the mental processes that enable sound judgement under changing conditions. Exhibit 1 sets out the cognitive skills that training builds and that the QF32 crew brought into the cockpit.

EXHIBIT 1

Pilot training is systematic cognitive development: seven skills that make information management — and therefore mutuality — possible.

The cognitive skills observed in commercial pilot training and how each is developed (Jenkins, from professional observation).

Cognitive skill	How pilot training develops it — or what capability the pilot must bring
Conceptualisation	Pilots build mental models of aerodynamics, weather, navigation, aircraft systems and human performance. They learn to understand principles rather than memorise procedures.
Orientation	Constant awareness of position, direction, altitude, time, terrain, weather, fuel state and aircraft configuration builds situational awareness that is continually updated as conditions change.
Attention and focus	Selective attention — knowing what deserves attention now while monitoring everything else. Scan patterns for instruments and the outside environment become disciplined habits.
Analysis	Every abnormal indication requires interpretation. Pilots analyse weather, performance, system failures, fuel and risk before deciding on a course of action.
Strategy and process	Flights are planned strategically before departure. In flight, pilots follow structured processes — checklists, standard operating procedures, decision-making models — while adapting them to changing circumstances.
Sequencing	Aviation is highly sequential: pre-flight, engine start, taxi, take-off, climb, cruise, descent, approach, landing, shutdown. Performing steps out of sequence increases risk; sequence supports accurate measurement and provides behavioural logic.
Rule recognition	Pilots learn regulations, aircraft limitations, airspace, weather minima and emergency procedures — and, more importantly, *why* the rules exist and when they apply.

These are the foundations for all information management — sensory, visual, auditory and linguistic — which in turn drives behaviour. Without them, in Jenkins's words, there is little chance of developing mutual relationships or even flying a plane. Source: Jenkins (2026), this paper.

Cognitive development strengthens recognition, categorisation, sequencing, comparison and pattern detection, enabling a person to organise information into meaningful relationships and discover the logic beneath it. On this view intelligence is the capacity to acquire, organise, interpret, communicate and apply information: to transform information into models of understanding and to convert those models into effective processes, decisions and adaptive action. It is not the possession of knowledge but the capability to manage information in ways that improve collective learning, comprehension, reasoning, communication and adaptation. Jenkins proposes a compact extension of the first paper's theory: intelligence is the capacity to transform information into adaptive understanding, coordinated action and mutual capability.

The ladder of cognition

A second feature of pilot training is that it progresses through increasing levels of cognition, and the progression is itself a description of how a person, or a crew, comes to judgement (Exhibit 2). The first rung is the hardest and the most easily skipped: collecting data dispassionately, without the temptation to interpret it, without assumptions and bias. Only then come patterns, the logic in the patterns, the facts that can be accepted as knowledge, and the understanding of why those facts hold. Application, integration and adaptation follow, and the top of the ladder is judgement: rational decisions under uncertainty, made on the basis of underpinning principles and practices rather than rules that no longer fit.

EXHIBIT 2

Judgement is the top of a ladder that begins with collecting data without interpretation; each rung depends on the one below.

The ten levels of cognition through which pilot training progresses (Jenkins, from professional observation).

Level	What it is
1 · Collecting data dispassionately	Looking at data without assumptions, bias or the temptation to interpret it.
2 · Looking for patterns	Relationships that recur — when this appears, that also appears consistently.
3 · Looking for the logic in the patterns	Does the pattern follow an "if this, then that" logic?
4 · Identifying consistently relatable data as fact	Accepting information as knowledge when it keeps recurring.
5 · Knowledge	Learning facts, rules and procedures.
6 · Understanding	Knowing why the rules exist and the logic of what has been established as fact.
7 · Application	Performing procedures strategically, with reason and logic, and with a clear understanding of risk.
8 · Integration	Coordinating multiple systems simultaneously.
9 · Adaptation	Modifying actions when conditions change unexpectedly; deciding under pressure by following the steps above.
10 · Judgement	Making rational decisions under uncertainty on the basis of principles and practices.

The ladder is the working definition of the "cognitive baseline" used throughout this paper, and it reappears in Section 7 as the thing that scaffolded AI assistance protects and unscaffolded advice erodes. Source: Jenkins (2026), this paper.

A teachable baseline. The claim that the baseline can be raised is not a hope. Jenkins's Essential Skills Profile, developed with Helga Rowe, measures applied capability rather than acquired knowledge — how effectively a person can perceive, organise, reason about and act upon information — and has been used for almost three decades to identify cognitive roadblocks and to track how they change after intervention. Its origin was a formative observation of children who had been deprived of social interaction from birth and had not reached expected developmental milestones, which challenged the assumption that cognitive function develops innately with maturation. Cognitive skill, on the evidence, is a renewable, teachable process — a conclusion Feuerstein (1980) reached from the other side of the world, and one on which everything else in this paper depends.

3 · THE CHAIN

Weak information management ends in lost trust — and the record shows the chain at organisational scale

Mutuality depends, in large part, on the ability of individuals to manage information effectively. When people lack the cognitive skills to observe accurately, organise information, distinguish the relevant from the irrelevant, interpret context and communicate clearly, misunderstanding becomes more likely. Persistent misunderstanding increases uncertainty; uncertainty reduces trust; as trust declines, mutual relationships weaken and collaboration and collective adaptation become harder. Jenkins sets the chain out in six steps.

1. **Weak cognitive information management** → difficulty processing and interpreting information.
2. **Poor interpretation** → miscommunication through poor comprehension.
3. **Misunderstanding** → uncertainty about another person's intentions, competence or meaning.
4. **Uncertainty** → reduced trust.
5. **Reduced trust** → weaker mutuality and poorer collaboration.
6. **Weaker collaboration** → diminished Human Adaptive Intelligence.

The relationship also runs in the opposite direction. Greater trust encourages people to share more information openly; more complete information improves understanding; better understanding strengthens trust and mutuality. The result is a positive feedback loop that continually improves the environment of mutual trust in which Human Adaptive Intelligence flourishes. Misunderstanding and mistrust produce the negative loop: division, alienation, adversarial confrontation and a desire to parade superior qualities that closes the mind to intellectual growth.

The same chain, drawn as a whale

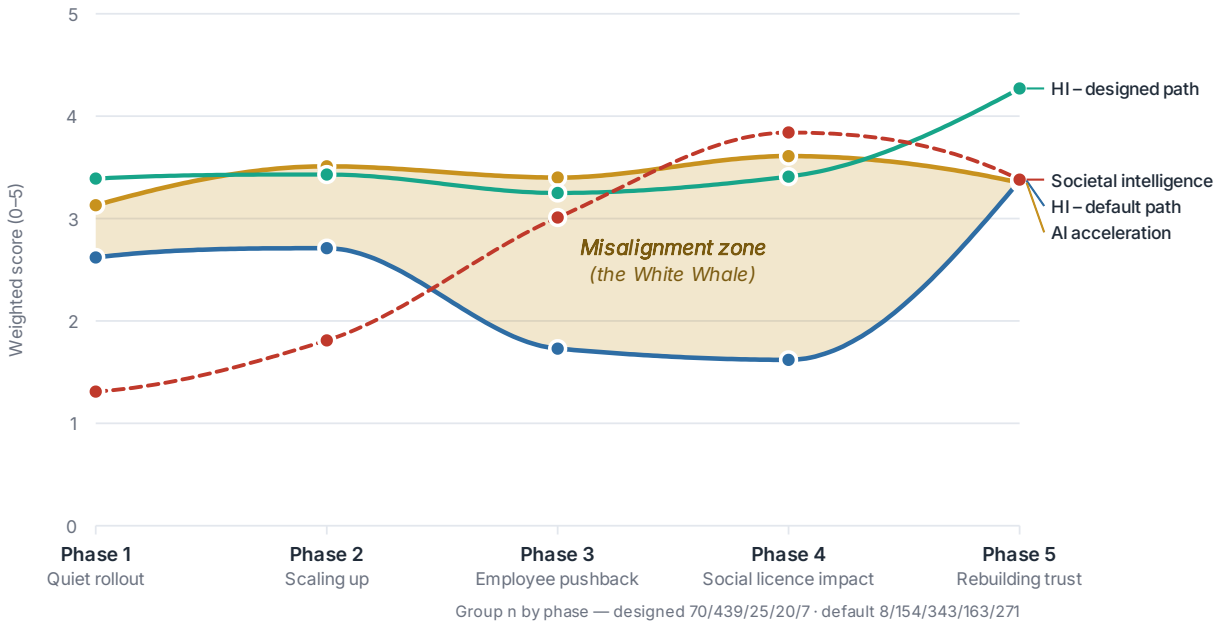
That is a description of one person, or one crew. Since 2025 the Fifth Revolution research programme at Building Mutuality has been watching something adjacent at a different scale: a source-referenced evidence bank of documented AI transitions — organisations, institutions, statutes, court decisions, labour agreements and evidence studies — each coded on five 0–5 scales and placed in one of five phases. It reached 1,500 cases in September 2026 and is published as *The White Whale Report 2026* (Byrne, Handley and Klaassen, 2026). The bank was not built to test Jenkins's theory, which is precisely why it is useful here: if the chain is real, it should be visible in a record assembled by people who were

not looking for it. Exhibit 3 shows what the record looks like when four of its lines are drawn across the five phases of an AI transition.

EXHIBIT 3

The White Whale at N=1,500: AI acceleration runs ahead of default-path human capacity through the contested phases, and the designed path does not take the dive.

Weighted phase means, 0–5 scale, five phases, N = 1,500 quantitative cases; designed and default are mutually exclusive coded groups (group sizes below the chart).



Lines interpolate five cross-sectional phase means and are conceptual, not longitudinal trajectories. Designed path = Alignment-outcome cases in Phases 1, 2 and 5 plus the 45-case contested-phase designed set (HI transition capacity ≥ 3.0 in Phases 3–4); default path = the remainder. The shaded body is a modelled gap between coded groups, not a causal estimate. Source: Fifth Revolution evidence bank, N=1,500 quantitative cases (Frozen Package N1500 v1, September 2026); evidence-tier weighted (A=1.0, B=0.75, C=0.5). Reproduced from *The White Whale Report 2026*, Exhibit 1.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

Follow two of the lines first. Gold is AI acceleration, the engine. Blue is human transition capacity on the default path, the driver, when no one designs for it. In the design window, when AI is quietly piloted and then scaled, the gap between them is narrow and beginning to open. Through the contested phases — when people inside start asking whether this is fair, and the issue then moves outside to courts, regulators, unions and the press — the engine keeps climbing while the driver collapses: human capacity falls to 1.73 and 1.62 while AI acceleration runs at 3.40 and 3.61. At settlement the lines close again. It looks like a whale. Read it the way the whale swims: tail at Phase 1, body through Phases 2 to 4, head at Phase 5, where trust is rebuilt.

The body of the whale is Jenkins's chain running through a whole organisation at once. People are handed a system whose outputs they cannot interpret, question or explain to the person it affects; misunderstanding multiplies; uncertainty about intentions and competence follows; and trust, which the bank measures as legitimacy, falls to 1.91 and 1.71 through the contested phases. Two more lines complete the picture. Red is societal intelligence — courts, regulators, unions, media and publics — rising as the gap widens. Green is human transition capacity on the designed path: the cases where people were enabled before the whale formed around them. The practitioner's reading is simple. AI is the engine; human intelligence is the driver; and every crash in the record happened where the engine outgrew the driver.

Measure the distance, not the adoption rate

The association is not loose. Across 1,500 cases the correlation between the size of the gap and the severity of what followed is r = 0.82, positive within every phase. Sort the cases by how far the engine runs ahead of the driver and the pattern is a staircase (Exhibit 4). Where people are ahead or level, no case in the record ends in backlash. Half a point behind, 3% do. One point, 11%. One and a half, 53%. Two points, 81%; beyond that, every case. Mean consequence intensity climbs the same stairs, from 0.88 to 3.41, and legitimacy falls down them, from 3.85 to 1.32. Adoption rate tells a leadership team how fast the engine is running. The number that tracks risk is the distance between the engine and the driver, and most dashboards carry the first number and not the second.

EXHIBIT 4

Backlash is absent where people keep pace and universal where AI runs two and a half points or more ahead; consequence intensity climbs the same stairs while legitimacy falls down them.

Share of cases ending in backlash, %, mean consequence intensity and mean legitimacy (0–5) by AI–HI gap band, N = 1,500.

Gap band (AI minus HI)	Cases	Ending in backlash	Mean consequence intensity	Mean legitimacy
HI ahead (uplift)	305	0%	0.88	3.85
Level	250	0%	1.08	3.61
+0.5	269	3%	1.22	3.32
+1.0	212	11%	1.72	2.86
+1.5	184	53%	2.23	2.13
+2.0	161	81%	2.79	1.66
≥ +2.5	119	100%	3.41	1.32

Gap = AI acceleration minus HI transition capacity, coded in half points; bands group the coded gaps ("+1.0" = gaps above 0.5 up to 1.0). Backlash is the coded outcome path. The staircase is a descriptive association between coded variables, not a threshold model; weighted r between gap and consequence intensity = 0.824, de-phased 0.678, within phases 0.67–0.81. Source: Fifth Revolution evidence bank, Frozen Package N1500 v1; *The White Whale Report 2026*, Exhibit 3.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

Same technology, different design

The clearest reading comes from the 305 cases, one in five, in which the people were enabled ahead of the machine: human capacity scored above AI acceleration. The bank calls this uplift, and QF32 is an uplift case — the most prepared aviator in the cockpit reached the edge of his own cognition, and the crew around him were capable, briefed, trusted and awaited. In the 305 uplift cases, 58.4% reached alignment, people and AI advancing together, and a further 40.7% reached a negotiated or legislated settlement. Not one ended in backlash. Not one reached a severe consequence. In the other 1,195 cases, 31.8% ended in backlash and 19.4% reached a severe consequence; legitimacy averages 2.70 against 3.85, and consequence intensity is more than twice as high (1.87 against 0.88). The most effective AI risk control the bank can find is the capability and trust of an organisation's own people — which is to say, the pillars of Section 2.

What the shape can and cannot say. The bank is a purposively assembled record, not a random sample; its scores are interpretive codes drawn from the public record, independently double-coded on a blinded sample; and every contrast above is a descriptive difference between coded groups, not a causal estimate. Across seven builds, seven regions and every batch of entry tested, the shape has held, and it holds on the strongest evidence alone. That is enough to say something the QF32 crew would recognise: the crash happens where the engine outgrows the driver, and the size of the gap, not the existence of the tool, is the risk.

4 · SHARED COGNITION

Crew Resource Management is the green line's playbook — and the order matters

Where CRM came from

CRM grew out of a NASA workshop in 1979 that reviewed a run of airline accidents and found that the crews had not lacked technical skill; they had failed at communication, decision-making and leadership, and more than once a captain's error had gone unchallenged by the people beside him (Helmreich, Merritt and Wilhelm, 1999). Aviation's response was not better pilots. It was a different relationship between pilots.

It became increasingly obvious to Jenkins, observing pilots at work, that modern aviation is built on cooperation rather than individual brilliance. Those in aviation recognise that safety emerges from shared cognition rather than isolated expertise: a captain and first officer live in an environment of shared minds, so that each can anticipate the other's erroneous thinking. Crew Resource Management teaches pilots six disciplines: to communicate clearly, using exact words that convey precise meaning; to share mental models and make sure that another person's model aligns with one's own; to challenge assumptions respectfully; to cross-check each other's thinking, repeating back what has been understood; to distribute workload so that no one is overloaded; and to build trust through competence and reliability — including the transparency to own up to mistakes, and the willingness to ask a colleague "are you OK?".

Read the case notes of the 305 organisations that kept people ahead of the machine and three moves recur so consistently that they amount to a playbook. **Change is negotiated rather than announced**: works councils, unions, guilds, co-design, a signed agreement before the cranes arrive. **Governance is in place before scale**: an ethics board with the authority to pause, an audit invited before it is demanded, a human review that everyone affected knows exists. **Capability is built at the point of rollout, not after it**: an academy, a certification, a learning entitlement, two years of study before anyone touches the line. Cases carrying each signal show uplift at about twice the rate of cases without it, and among the 45 that held capability through open conflict the three signals are four to six times as common as among the 506 that collapsed (Exhibit 5).

EXHIBIT 5

The three moves of the green line are Crew Resource Management translated out of the cockpit, and each doubles the uplift rate where it is present.

Share of cases with HI transition capacity above AI acceleration, %, where each signal is present in the case record against where it is absent, N = 1,500; the CRM discipline each move expresses and the keystone it belongs to.

Green-line move	Uplift where present	Uplift where absent	The CRM discipline it expresses	Keystone
Change negotiated, not announced	43.8%	18.9%	Share mental models; challenge assumptions respectfully	Mutuality
Governance in place before scale	35.7%	16.8%	Cross-check each other's thinking; distribute workload	Culture, Integration
Capability built at rollout	38.2%	18.6%	Build trust through competence and reliability	Capability

Signals are keyword-coded from case names and rationale notes (negotiated: union, works council, co-design, agreement…; governance: oversight, ethics, audit, contestability, human review…; capability: reskill, academy, training, certification…); hypothesis-grade and sharing their source with the scores they are compared against. Source: Fifth Revolution evidence bank, Frozen Package N1500 v1; *The White Whale Report 2026*, Exhibit 5.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

Negotiated change is the shared mental model and the standing invitation to challenge. Governance before scale is the cross-check, institutionalised: someone whose job is to ask "are you sure?" and be thanked for it. Capability at rollout is trust built on competence and reliability, which no one can extend to a colleague, or a system, they do not understand. The sixth discipline, the plain human question *are you OK?*, is the one the record cannot code and the one that Section 6 argues matters most.

The order of operations

What the record adds to CRM is the order. *The 5th Revolution* sets out four keystones of healthy integration between human and artificial intelligence — Mutuality, Culture, Capability and Integration — and argues that they are a sequence rather than a list: an arch is not begun with the top stone. Most organisations do exactly that. They buy the tool, integrate it and turn to the people afterwards, with training as an afterthought and governance when the regulator asks. The bank records how that order ends. In the 945 cases where AI acceleration ran ahead of human capability, 22.2% reached alignment and 40.2% ended in backlash. In the 555 where human capability kept pace with or led deployment, 56.2% reached alignment and none ended in backlash; legitimacy 3.74 against 2.45; severe consequences none against one in four. The 45 that held through open conflict reached a severe consequence in 8.9% of cases against 43.3% for the default group. Human intelligence first is not the slow option. In this record it is the one that arrives.

The cases are ordinary organisations doing recognisable things. IKEA retrained 8,500 call-centre agents as interior advisers when its bot took the routine calls: redeployment designed before displacement (N1000-001). Hitachi deployed generative AI through union consultation and an internal certification (N1000-040). Lloyds Banking Group built an AI academy with a whole-workforce literacy target after a 7,000-person trial and before the agentic wave (N1500-035). Aftonbladet's journalists co-designed the newsroom's AI tools and kept sign-off (N1500-150). None of them avoided the gap between the engine and the driver: in the contested phases AI runs ahead in 60% of the designed cases too, by 0.36 on average and never by more than a point and a half, against 1.76 in the default group, where it is ahead in every case. The size of the gap, not its existence, separates the paths.

The desert is the enterprise. Uplift appears in 8.9% of large-enterprise cases against 26.4% of government and public institutions and 48.8% of research bodies; even inside the design window the enterprise rate is 10% against 32%. The organisations best resourced to build human capability alongside AI are, so far, the least likely to have done it. Klaassen's explanation, from two decades inside enterprise transformation programmes, is structural: the platform budget and the capability budget do not arrive at the board table as one decision. The first comes with a vendor, a date and a business case; the second arrives later, if at all, as a people cost. Nobody sets out to leave people behind. It is what happens when two decisions that belong together are made separately — a governance structure, in Jenkins's terms, with the feedback loop removed. Uplift is a design problem, not a money problem.

5 · CONNECTION AND ACCOUNTABILITY

Human Adaptive Intelligence flourishes when the person doing the work can change the decision

The proposition

Human Adaptive Intelligence flourishes when the person performing a function is cognitively connected to observing, evaluating and improving that function. It diminishes when observation, decision-making and process improvement become separated from the people experiencing the process. — Jenkins (2026)

The shared-intelligence culture of the cockpit was once shared by the engineers who built the aircraft. Jenkins traces what happened when, at Boeing, it was not. The company had used a relatively flat, integrated governance model in which workers, engineers and corporate management collectively established the principles and processes for manufacturing, and communication and decision-making were transparent and guided by established quality-control rules. Under new management the governance of its divisions was separated; bottom-up engagement between engineers and management gave way to a hierarchical, top-down structure with restricted feedback loops; collective intelligence fragmented and political factions took control. Decisions were made without input from the people experiencing the process. The result was disastrous, with accidents and loss of life.

The congressional investigation into the 737 MAX reads like a case written to that specification. After eighteen months, the House Committee on Transportation and Infrastructure (2020) described production pressures that overrode engineering judgement, faulty assumptions about how pilots would respond to a system many of them had not been told existed, a culture of concealment, and an oversight arrangement in which the manufacturer had come to grade much of its own work. Every one of those is a severed connection between the people who could see the problem and the people who could change the decision. The engineers had not lost their cognition. They had lost their standing.

Accountability is a relationship, and it has to be located

The proposition turns out to be testable, because the most confused conversation in the AI story is the one about accountability. Regulators have answered the "responsibility gap" (Matthias, 2004) with people: Article 14 of the EU AI Act wants natural persons able to "disregard, override or reverse" a high-risk system, and Australia's voluntary AI safety standard opens by saying leaders cannot delegate or outsource accountability. The critics have shown why a human in the loop so often fails; Elish (2019) named the "moral crumple zone", the nearest operator who absorbs the blame for a system they could not control. The courts have been less confused: told by Air Canada that its chatbot was a separate legal entity, a tribunal replied that it "should be obvious" the airline was responsible for its own website (*Moffatt v Air Canada*, 2024). Davies (2024) supplies the missing piece in one line: the extent to which

you are able to change a decision is precisely the extent to which you can be accountable for it. Read beside Jenkins's proposition, the two are the same sentence. Cognitive connection to a function *is* the ability to change it.

The evidence bank read 194 case records and coded where accountability for the AI-affected outcome actually sat (Exhibit 6). In the 69 enterprise and small-business cases where people were enabled ahead of the machine, accountability was located by design in 84% and displaced in none. In the 45 cases that held human capacity through open conflict — the hardest test in the bank — 91%. In a random sample of 80 backlash cases, none: 70% had displaced it, onto the system, onto the people affected or by disclaiming it, and the rest record a court, a regulator or a campaign locating it afterwards, at the organisation's expense. Robodebt, in which automated debts were raised against welfare recipients with no human able or willing to stop the scheme, is the displaced case in its purest form (Royal Commission into the Robodebt Scheme, 2023).

EXHIBIT 6

Where the green line locates accountability, the whale displaces it: 84% of enterprise uplift cases and 91% of the contested-phase designed set locate it by design; none of the backlash sample does.

Where accountability for the AI-affected outcome sat, % of cases coded from the case record: enterprise and SME uplift (n=69), contested-phase designed set (n=45), seeded random backlash sample (n=80), N = 1,500 bank.

Group	Located by design	Displaced or disclaimed	Imposed from outside after harm
Enterprise and SME uplift cases (n=69)	84%	0%	—
Contested-phase designed set (n=45)	91%	—	—
Random backlash sample (n=80)	0%	70%	30%

Coded by the lead author of the report from case names and rationale notes against a released rubric, and independently double-coded by the external validators. Displaced onto the system: e.g. Robodebt's automated debts; onto the people affected: gig workers, welfare recipients; disclaimed: Air Canada's chatbot; imposed = accountability arrived from outside after harm. Descriptive throughout. Source: Fifth Revolution evidence bank, Frozen Package N1500 v1; *The White Whale Report 2026*, Exhibit 6 and Appendix C.8.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

Read the cases that located accountability and it sits in one of four places. **In a role with authority**: a named person or team who decides, can override and can stop, and everyone affected knows it — Kaiser Permanente's early-warning algorithm in hospital wards is answered by a nurse-led response team; the machine proposes, the clinician decides (N1500-068); Major League Baseball let the automated ball–strike system hear the appeal and kept the umpire's call (N1500-062). **In an agreement with the people affected**: Hitachi's union consultation and certification (N1000-040); Te Hiku Media's Māori speech AI under community ownership and indigenous data sovereignty (N1500-074). **In an institution that can pause or audit**: Axon's own ethics board paused facial recognition on police body cameras, in the phase where most organisations cannot stop anything (N760-692). **In capability that keeps people able to hold it**: DBS upskilled more than ten thousand staff while it scaled enterprise AI (N1000-041); Daikin pays its graduates for two full years of AI study before they touch the factory (N1500-125). A person who cannot explain a decision cannot be accountable for it, which is why the fourth place is not an afterthought to the other three. It is the cognitive baseline of Section 2, funded.

The bank also holds twenty-nine reversals and path switches, and they are one event repeated. Klarna rehired human service agents after cost had become, in its chief executive's words, "a too predominant evaluation factor" (N1000-017); the Commonwealth Bank reinstated forty-five roles after a union challenge showed the work had not gone away; Ford rehired engineers for judgement work. In each

case accountability moved back to people after it had been placed on a system, at a price that would have been far lower in the design window. These are the QF32 cockpit in reverse: the connected, mutually engaged process rebuilt after the hierarchy had tried to do without it.

> **Flatter, or connected?** Jenkins argues for flatter management models, mutually engaged with people at the coal face of the operation. The record refines the claim. The green-line organisations are not necessarily flat; they are connected: the person performing the function can observe, evaluate, contest and improve it, and the people above them treat that as the system working rather than as insubordination. It is what the two check captains did when they set rank aside. The bank's word for it is *located accountability*. Jenkins's word for it is *mutuality*. They describe the same cockpit.

6 · EMOTIONAL ENERGY, MUTUALITY AND MATTERING

Trust is the residue of being noticed, valued and relied upon — and it has to be banked before the unknown arrives

A person's neurological energy, Jenkins observes, can be directed toward beneficial or non-beneficial outcomes. Unchecked, emotional energy emerges as anxiety, fear, panic, terror or freeze; directed toward a motivating goal, it produces beneficial outcomes. In a crisis, pilots have clear paths for enquiry, investigation and system analysis, and clear paths of appropriate communication, and their energy is channelled down them rather than into panic. This is Human Adaptive Intelligence's understanding of emotional intelligence: an individual or group collaboratively chooses how and where to place emotional energy. De Crespigny, about to lose the capacity to hold the aircraft in his mind, chose the path that was best for everyone — delegating the collection and interpretation of the warning codes and giving everyone else a clear path for data collection and interpretation. People become adaptive when they control their emotional energy, and the sequence Jenkins draws from pilot training follows from it: cognitive development leads to better information management; better information management to clear and transparent communication; clear communication to mutual comprehension; and mutual comprehension to trust. When those conditions are perceived to be in place, Human Adaptive Intelligence begins to flourish.

Mutuality: from the book to the bank

Mutuality: The Future of Trust (Byrne and Jenkins, 2024) described mutuality as a tool, a process and a practice that builds respect and trust — outcomes organisations had long sought and found elusive — and drew the same distinction between endorphin-based environments, where energy is directed positively and gradually into ownership and discovery, and adrenaline-based ones, where adversarial management redirects it into anxiety, self-protection and politics. Each point of mutuality found adds a thread; the threads weave into a string; the string becomes a rope that can withstand robust debate, disagreement and the occasional breach of group rules. The book's central practical claim was that the most efficient way to make mutuality possible is through cognitive development, because without it people may not understand what their mutual responsibilities to one another are. Sections 2 and 3 of this paper supply the mechanism behind that claim.

The 5th Revolution gives mutuality a sharper operational definition, because the AI era tests it in a specific way: a power-sharing practice in which the people who will use, be affected by or carry the consequences of a system are treated as co-authors rather than users, realised through voice, through

explainability, through AI that people can interrogate and influence, and through the willingness to treat challenge as information rather than insubordination. The question that tests it is simple: is the organisation designing with people, or only for them? The book opens on a case Jenkins's chain explains exactly. A regional council introduced a generative AI tool to assess planning applications; it did in seconds what had taken officers hours, and within months it sat largely idle. Staff had reverted to manual work, and senior officers quietly overrode most of its recommendations after hours. There had been no orientation, no training, no feedback cycle and no shared conversation about what human judgement would now mean. The system worked; the people did not. Those officers were not resisting innovation. They were resisting exclusion: the information they needed to interpret the tool had never been given to them, so uncertainty was the only rational state, and trust could not form.

The trust bank has a clock on it

The evidence bank sees Jenkins's positive loop, and it sees when the loop has to be running. Its nearest measure of trust is legitimacy — whether the people affected found the way AI was introduced fair and worth their trust. In the design window, cases where legitimacy already scored 3.5 or higher show uplift in 35.1% of cases; below that line, 2.6%; at 4.0 or higher, 52.7% (Exhibit 7). Through the contested phases the default group's legitimacy falls to 1.91 and 1.71 and does not recover until settlement. Trust is the residue of repeated experiences of being noticed, valued and relied upon, and it cannot be created quickly under pressure. An empty trust bank at Phase 3 cannot be filled in Phase 4. This is the positive loop with a clock attached: it has to be running before the unknown arrives. On QF32 it was running before the engine failed, because the captain had briefed the crew not to stay quiet and not to defer simply because a check was under way.

EXHIBIT 7

Legitimacy in the design window is the strongest early signal in the bank — uplift in 35% of Phase 1–2 cases scoring 3.5 or higher, 3% below it — and the default path's legitimacy collapses through the contested phases.

Share of Phase 1–2 cases with HI above AI, %, by legitimacy band (n=671); mean legitimacy by phase, uplift cases against the rest, 0–5 scale, N = 1,500.

Design-window legitimacy band (Phases 1–2, n=671)	Share of cases showing uplift
Below 3.5	2.6%
3.5 or higher	35.1%
4.0 or higher	52.7%
Default-path legitimacy by phase	**Phase 3: 1.91 · Phase 4: 1.71 · Phase 5: settlement**

Legitimacy is a coded 0–5 score of whether the change was accepted as fair and trustworthy by those affected. Phases are coded at the most recent observable position, so the design-window association is not a longitudinal prediction; HI transition capacity and legitimacy correlate at 0.95 and share a coding source. Source: Fifth Revolution evidence bank, Frozen Package N1500 v1; *The White Whale Report 2026*, Exhibit 7.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

This is why AI is best understood not as a magic answer or a rival mind but as an amplifier. It takes the culture it lands in and turns up the volume: in high-trust organisations it behaves like a multiplier — faster learning, better decisions, more room for human connection — and in low-trust organisations the same technology behaves like an accelerant — faster blame, faster surveillance, faster replication of bias (Byrne, 2026b). Hitachi and Klarna deployed comparable technology into customer-facing work within a year of each other; one did it through its union with a certification, the other automated its service before reading its customers' logic, and then reversed (N1000-040; N1000-017). The machine

did not decide which of those happened. The relationship around the machine decided, and the machine made it louder.

Mattering: the question beneath the pillars

Jenkins's positive loop begins with a behaviour: people share information openly. Everything downstream depends on it, so it is worth asking what, in a human being, opens or closes that door — because information is not withheld at random. It is withheld from people and systems that have signalled it will not be received. The answer developed in the article that accompanies the White Whale Report (Byrne, 2026d) is mattering. Rosenberg and McCullough (1981) defined it as the felt sense that others are aware of you, attach importance to you and depend on you. Prilleltensky (2020) gives it two sides, feeling valued and adding value, and finds that fairness produces wellbeing largely through it. Flett (2018) shows that its absence is not a low score but a distinct state, anti-mattering, that does more harm than mattering does good. It rests on the need to belong, which Baumeister and Leary (1995) showed to be as basic a drive as hunger. Mercurio (2025) reduces the practice of conferring it to three verbs: noticing, affirming, needing. The sixth discipline of CRM — the willingness to ask a colleague "are you OK?" — is noticing in its purest form.

Read QF32 through mattering and the cockpit becomes clearer still. The first officer's reading of the fault messages was awaited. The second officer's calculations were relied upon. The check captains could set rank aside because in that cockpit rank was not the source of significance; contribution was. Everyone mattered, and so everyone gave what they had. Jenkins is right that a minimum cognitive baseline is required for mutuality. Mattering decides whether the cognition people have is offered or kept back.

Now consider the far more common case in the record: AI arriving not as a shared instrument but as something deployed over people's heads. The technology has not changed. What has collapsed is the signal that anyone is awaited or relied upon. Chun, De Cremer, Oh and Kim (2024) found that the main harm of algorithmic evaluation is not perceived bias but perceived disrespect: people conclude that their particular characteristics were never considered, that they were processed rather than seen. The American Psychological Association's 2023 survey of 2,515 workers found that 38% worried AI would make some or all of their duties obsolete, and that the worried were nearly twice as likely to feel they did not matter to their employer (41% against 23%) or to their colleagues (37% against 17%). To be trusted with a decision is the clearest evidence a person can receive that they are still needed. To be made the crumple zone for a decision they could not change is anti-mattering in its purest form. Some of the most capable people in the record — Boeing's engineers, the council's officers, Klarna's service agents — withheld their intelligence not because their information management was weak but because they had been told, in the most credible currency there is, that they no longer counted. Hodson (2001) named the specific insult: pseudo-participation, being consulted but not heard.

> **Mattering can be designed.** The three green-line moves are what it looks like when it is. Negotiated change says *your voice counts*. Visible governance says *your judgement is still the point*. Capability built early says *you belong in the future being built*. Read psychologically, the sequence of mutuality, culture, capability and integration is an architecture of significance. The bank cannot yet measure mattering directly — no case in it administered a validated instrument, and that study is first on the research agenda (Section 10). But the pattern was there before the word for it, and the word tells leaders where to look: a legitimacy score cannot be fixed, but people can be noticed, the specific difference they make affirmed, and their being needed shown, in the daily moments AI has not automated and in the systems it has.

AI that returns process intelligence to the person lifts the least experienced most; AI that removes it widens the gap

Two questions

Automation asks how a process can be made to require less human intervention. Human Adaptive Intelligence asks how the human participant can be made better able to perceive, understand and improve the process through the use of AI. Those, Jenkins writes, are very different objectives.

Captain de Crespigny held that people must always remain in control of systems, because only people can make the dynamic trade-off decisions that cannot be programmed into a computer; tasks, technology and environments must therefore be designed for people, and people must have the competency to handle the situations that can be foreseen and those that cannot. Jenkins draws the consequence for AI. If AI simply tells a worker what to do, it further separates the worker from cognition. If it helps the worker observe, compare, question, model, test, refine, understand and improve, it returns process intelligence to the person doing the work. AI could be either destructive or enormously beneficial to Human Adaptive Intelligence, and the difference is not in the model. It is in the direction the model is pointed.

The field evidence of the last three years shows that difference with unusual clarity, because it has been running both experiments at once. The landmark studies of 2023 to 2025 entered the evidence bank this year as scored cases, and read together they describe one mechanism (Exhibit 8). In a customer-support centre, generative assistance raised productivity most for the least experienced agents, by around a third, while doing little for the most experienced (Brynjolfsson, Li and Raymond, 2025). In professional writing tasks the lowest-graded writers gained most and the spread narrowed (Noy and Zhang, 2023). Among consultants working inside the tool's competence, below-average performers gained 43% and above-average 17%; on a task that looked similar but lay outside it, those who leaned on the tool did worse (Dell'Acqua et al., 2023). The lowest-rated tutors in a large tutoring programme gained most when a copilot suggested expert moves in real time (Wang et al., 2024). And in two studies the effect ran the other way: Kenyan entrepreneurs already performing well gained from AI advice while those performing poorly lost ground (Otis et al., 2023), and experienced open-source developers completed real tasks 19% more slowly with AI tools while believing themselves about 20% faster (Becker et al., 2025).

EXHIBIT 8

Designed assistance lifts the least experienced most; unscaffolded advice widens the gap. Seven studies, one pattern.

Selected evidence studies held in the bank as scored cases; results as reported by the studies, populations and effect sizes their own.

Study	Setting	How "who gained" was measured	Direction
Brynjolfsson, Li and Raymond (2025)	Customer support, generative assistant	Tenure and prior skill	Least experienced gained most (about +34%)
Noy and Zhang (2023)	Professional writing tasks	Grade on a first, unassisted piece	Lowest gained most; spread narrowed
Dell'Acqua et al. (2023)	Consultants, inside and outside the "jagged frontier"	Baseline assessment task	Below-average +43%, above-average +17% inside the frontier; worse outside it
Wang et al. (2024)	Tutors with a real-time copilot	Tutors' prior ratings	Lowest-rated gained most (+9 points)
Otis et al. (2023)	Kenyan entrepreneurs receiving AI advice	Baseline business performance	High performers gained; low performers lost
Becker et al. (2025), METR	Experienced open-source developers	Experience with the codebase	19% slower with AI while believing themselves faster
Humlum and Vestergaard (2025)	Danish workers, population register	No baseline split	No measurable earnings or hours effect

"Least experienced" is a synthesis label: the four studies in which the people who started furthest behind gained most measured that starting point in four different ways (tenure, a first unassisted attempt, a baseline test, a manager's rating). What they share is a mechanism, not a population. Source: the studies cited; *The White Whale Report 2026*, Exhibit 12 and Section 5.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

What separates the studies in which the people furthest behind gained most from the two in which they lost is not the model. It is whether the assistance was built around a person doing a well-specified job, with someone accountable for how it landed, or dropped on people as advice. Education has a fifty-year-old word for the first kind: scaffolding, the support that lets a learner do today, with help, what they will do alone tomorrow (Wood, Bruner and Ross, 1976; Vygotsky, 1978). Scaffolding is the ladder of Exhibit 2 with a hand on the rail. It keeps the person on the rungs — observing, comparing, questioning, testing — rather than lifting them past the rungs to the answer. Advice does the opposite: it supplies the answer and removes the reason to climb.

Bainbridge's irony, and the pace finding

Aviation learned this before anyone. Bainbridge's "Ironies of automation" (1983) observed that the more a system is automated, the more its safety depends on the skill of the human who must take over when it fails, and the more that skill atrophies through disuse. De Crespigny's philosophy is the working pilot's answer to Bainbridge: he studied the manuals and the aerodynamics of every aircraft he flew precisely because automation breeds complacency, and when the automation drowned him in messages the understanding was still there to fall back on. The evidence bank's version of the same irony is its pace finding. Cases deployed at frontier pace show uplift in fewer than one per cent of 308 cases, in every cohort, while at a scaled but humane pace the uplift share has held at around a quarter for a decade. The engine got faster and more organisations chose the fastest lane. Nobody has yet kept people ahead in it.

A quieter risk bears on the first rung of the ladder, collecting data dispassionately. Knowledge workers who trust generative AI more report putting less effort into critical thinking, shifting from doing the

thinking to checking it (Lee et al., 2025), and a 2026 study found that AI assistance raised people's performance while leaving their judgement of that performance no better calibrated: smarter, its authors said, but none the wiser (Fernandes et al., 2026). Jenkins's first paper listed overreliance and the loss of independent reasoning among the risks of AI. The evidence suggests the risk is real, and specific about where it strikes: at reflection, the rung between knowledge and understanding, which no tool can climb for a person.

> **Direction is a choice.** None of this is an argument against the tools. It is an argument about their direction, and economists have drawn the same line from the other side. Brynjolfsson (2022) calls the pursuit of human-like AI that substitutes for people rather than complementing them the Turing Trap; Acemoglu and Johnson (2023) show, across a thousand years of technology, that whether a machine raises or lowers the value of ordinary people's work is a choice, not a law. Jenkins's two questions are that choice, put to any organisation buying AI this year. The record says the answer is visible in whose judgement the tool amplifies — and that where it amplifies the judgement of the people who start furthest behind, the whole organisation gains.

8 · THE FOUR CAPABILITIES

What high human capacity is made of maps onto the cognitive foundations

Jenkins's table of pilot cognition describes what a mind must be able to do. When the case records behind the highest human-capacity scores in the bank were read for what the people in the green-line organisations were actually good at, four capabilities kept surfacing (Byrne, 2026c; Exhibit 9). None is a technical skill or a personality trait. Each is a trainable habit of reading reality where the model flattens it, and each rests on foundations Jenkins has already named. AI is good at maps. Human intelligence is what becomes visible when the weather changes faster than the map.

In the good cases somebody asked what the decision meant here, noticed what the pattern missed, read what the place really rewarded, and invented a better arrangement of the work rather than defending the old one. None of that lives in a model. All of it can be developed — and Jenkins's career is the evidence that the foundations beneath it can be developed too. Klaassen adds a layer under applied creativity that Jenkins's emphasis on humility anticipates: inventing the better arrangement of the work also needs the people already holding that ground to allow it. IKEA's agents could become interior advisers only because whoever held customer expertise before was willing to let the role move. In her experience the invention is rarely the hard part; the concession is. That willingness is what the check captains showed, and it is as teachable as the invention itself.

EXHIBIT 9

The four capabilities behind high human transition capacity are the cognitive foundations of pilot training, put to work in organisations with a powerful new tool in them.

The four capabilities read from the case records (Byrne, 2026c), the question each asks, and the cognitive foundations from Exhibit 1 on which each rests; cases in the record with identifiers.

Cognitive foundation (Exhibit 1)	Capability in organisational life	The question it asks	In the record
Conceptualisation; rule recognition (knowing *why* the rules exist)	**Context** — seeing the wider social, policy and institutional frame a choice sits in, so optimisation stays aligned with values, legitimacy and purpose	What does this mean here?	Te Hiku Media's Māori speech AI under indigenous data sovereignty (N1500-074); Bhashini, language AI as public infrastructure (N1000-139). Absent: the UK exam algorithm, technically defensible and publicly illegitimate (FR-193)
Orientation; attention and focus; analysis	**Circumstance** — reading the local situation, exceptions, stakeholders, timing and constraints, so a generic recommendation is never applied blind	What is true here that the pattern misses?	Kaiser's alerts answered by a nurse-led team (N1500-068); CottonAce pest alerts routed through lead farmers (N1500-113). Absent: Zillow, whose pricing model met a market it could not read (N760-700)
The CRM disciplines: communicate clearly, share mental models, challenge respectfully, cross-check	**Social logic** — the hidden operating system: what is safe to say, what is really rewarded, whose voice counts, whether explaining is part of the work	What does this system really reward, and whose voice counts?	Hitachi's union consultation and certification (N1000-040); the Culinary Union's AI provisions (N1000-013). Absent: Klarna, which automated its service before reading its customers' logic (N1000-017)
Strategy and process; sequencing; integration; adaptation	**Applied creativity** — inventing new arrangements of roles, rules and relationships so that people and the machine both fit; redesigning the work rather than bolting AI onto it	What could this work become?	IKEA's agents became interior advisers (N1000-001); Toyota put master craftsmen back on automated lines (N1500-134); Boliden's miners moved to surface control rooms (N1500-152). Absent: Henn na Hotel, which fired half its robots and rehired humans (N1500-128)

The mapping is interpretive and not one-to-one; it is offered because it says the four capabilities can be taught, since their foundations can be. The capabilities have since been coded independently across the uplift and designed cases by the external validators. Source: *The White Whale Report 2026*, Section 2 and Appendix D.5; Jenkins (2026), this paper.

Building Mutuality · Mutuality — The Future of Trust Project · White Paper 2

9 · THE TURN TO HUMAN INTELLIGENCE

In 2026 the people building the machines began describing the driver

For three years the argument about AI was entirely about the engine. In 2026 the people building the systems began to say what de Crespigny had said about automation: that people must remain in control.

On 12 September 2026 Dario Amodei asked the industry to pace the frontier, warning that AI "could outrun our ability to understand and control these systems"; Sam Altman agreed within hours and Elon Musk wrote "Dario is right". On 20 September Anthony Albanese called for a US–China framework on AI modelled on the nuclear non-proliferation treaty so that "humans remain in control". In June, Satya Nadella told the *Wall Street Journal* that AI companies "have to do the hard work in earning the social permission" of the public, and that leaders who see AI as a job-elimination tool are looking at the

technology wrong; the task is to reorganise the job around what people can do. Klarna's chief executive, having automated customer service and reversed, concluded that "really investing in the quality of the human support is the way of the future for us". McKinsey's *State of Organizations 2026* reports executives saying five dollars should be spent on people for every dollar spent on technology, and 86% of leaders saying their organisations are not prepared to adopt AI in day-to-day work.

Almost none of them used the words human intelligence. Every one of them described it: *our ability* to understand and control; *humans* in control; *social permission*; the quality of the *human* support. Side by side they point at the same missing thing — a capability inside people to understand, to judge, to stay in control and to earn permission. Guardrails restrain the machine. They do not build the human. Jenkins's reversed question from the first paper — not how should humans adapt to AI, but how should AI be developed to strengthen Human Adaptive Intelligence — has stopped being a minority position.

The more consequential shift is quieter: organisations are beginning to fund the human side the way they fund the machine. The Commonwealth Bank has committed about ninety million dollars over three years to a Future Workforce Program developed with its unions, aimed at the judgement, critical thinking, problem-solving and empathy its chief executive expects higher-value roles to need (Commonwealth Bank, 2026). Since February 2025 the EU AI Act has required a sufficient level of AI literacy in the people who operate AI systems. The caution is in the same record: Accenture trained 550,000 people in 2023 and in 2025 exited 11,000 "who cannot be reskilled" (N1000-005; N1500-023). Capability used as a filter is not the green line; it is the default path with a certificate attached.

> **Literacy is not human intelligence.** Most of what is being funded is literacy: necessary, and nowhere near sufficient. Literacy certifies a person's relationship to the tool. Human intelligence is a person's relationship to the decision, the context and the people it affects: the ladder of Exhibit 2, the four capabilities of Exhibit 9, accountability located in someone who can change the decision, and the order in which mutuality, culture and capability are built before integration. An organisation can certify its AI management system, its governance specialists and its people's tool fluency. It cannot certify its way to a cockpit in which everyone's contribution is awaited. That has to be built, and Sections 2 to 4 describe what it is built from.

10 · TOWARDS THE THIRD PAPER

Five questions the third paper must answer

Sections 1 to 4 named the pillars. Sections 5 to 9 showed that the record confirms them at scale, and confirms the order in which they have to be built. What neither has done is turn the proposition into instruments — things a leader, a teacher or a designer can pick up on Monday. That is the work of the third paper, and five questions from this one define it.

1. The Human Adaptive Intelligence test: how can AI that returns process intelligence to people be told from AI that removes it? Jenkins's cycle — observe, compare, question, model, test, refine, understand, improve — is more than a description. It is a design specification. The third paper should turn it into an evaluation standard that can be applied to any AI tool or deployment before it scales. Does the system show its reasoning in terms the person can interrogate? Does it leave the person able to compare and question, or does it hand over an answer? Can the person change the decision, and do they know they can? Does it build the person's model of the work or quietly replace it? The evidence bank offers the test cases already — Kaiser's nurses, Octopus's approvers and baseball's umpires on one side; Zillow, Robodebt and the crumple zones on the other — and the evidence studies of Exhibit 8 as the controlled version, scaffolding against advice. Norman's *Things That Make Us Smart* (1993) and

Brynjolfsson's Turing Trap frame the same question from design and economics. The prize is a standard that lets an organisation choose augmentation on purpose rather than discover automation by accident.

2. The cognitive baseline: what is the minimum, how is it measured, and who is at risk of being left below it? The claim that mutuality requires a minimum baseline of cognitive capability is the least comfortable and most important sentence in this paper, because it raises a question of inclusion. If working with AI well depends on the ladder from dispassionate data to judgement, then the AI era could open a developmental divide wider than the digital one. The evidence gives grounds for hope: assistance built around a person lifts those who start furthest behind, and the Essential Skills Profile's record shows that the foundations are teachable at any age. The third paper should say what the baseline consists of, how the profile measures it, how it changes under development, and what it means for the people the bank shows are most exposed — frontline, casual, older and less credentialed workers, and the students now entering a workforce that will assume this capability without teaching it. The HI Practitioner accreditation now in development is the test bed for whether it can be taught at scale.

3. Flatter organisations: what structure keeps the person doing the work connected to the decision? Jenkins argues for flatter management, mutually engaged at the coal face. The record refines the claim: the green-line organisations are connected rather than flat, and accountability lives in one of four places. The third paper should describe the organisational designs that keep the feedback loops intact as AI arrives — what Weick and Sutcliffe (2007) called deference to expertise in high-reliability organisations, the network of teams described in *Mutuality*, Ostrom's (1990) evidence that shared rules can be organised deliberately and durably, and Boeing as the standing warning of what happens when governance is separated from the people experiencing the process. Jenkins's own sketch of the model is the starting point: individual observation → relational understanding → experimentation → feedback → refinement → shared learning → mutuality → collective intelligence → improved function → new observation.

4. Education: what should be taught when prediction is cheap? The first paper argued that education should cultivate structured thinking tools rather than primarily facts. The printing press transformed only the societies that taught people to read; the AI era's literacy is not the ability to build the machine but the ability to read it, question it, put it to work and, when necessary, refuse it. The third paper should set out what that means for schools, for universities and for the executive classroom, drawing on the cases in the bank where a whole system tried — Kerala's programme for 80,000 teachers (N1500-110), Singapore's SkillsFuture pathways, Daikin's two years, the tutors and teachers of Exhibit 8. The ladder of cognition is the curriculum. The classroom is the cockpit, and the crew are the students.

5. Measurement: can the chain be tested as a leading indicator? The bank scores human capacity and legitimacy from the public record; it cannot see inside a person. The Essential Skills Profile can. A study that paired the profile with validated instruments for mattering and psychological safety, and tracked them through live AI deployments alongside the coded scores, would test this paper's central claim directly: that information management and trust move first, and that the whale can be seen forming before anyone pushes back. Two independent validators are re-coding a blinded sample of the bank now; the road to 2,000 cases is planned; and the first direct measurement of mattering is the next study on the agenda. The third paper should design that study.

THE INVITATION

Together is better — and the record now lets it be said with some precision

Underneath all five questions sits the conviction that runs through both authors' books and both papers: together is better. Not human intelligence alone, which reaches its limit in every cockpit and every crisis. Not artificial intelligence alone, which is only ever an amplifier of the culture that made it. But the two held in a relationship where the human remains the author, the machine remains the instrument, and the point of the arrangement is that more people, not fewer, get to think, contribute, belong and grow.

Jenkins found that relationship in a cockpit over Batam Island, and traced it back to the pillars beneath it: a cognitive baseline that can be raised, a chain that runs from information to trust and back, and a connection between the person doing the work and the decision that must never be cut. Byrne found it, drawn as a whale, in 1,500 organisations: the same chain at scale, the same order of construction, the same connection called by another name. It is one picture. The first paper proposed it. This paper shows that the evidence bears it out. The third will be about how to build it on purpose, for everyone, before the water gets deep.

Human intelligence is the main story. AI is the supporting one. Together is better.

ABOUT THE RESEARCH

How this paper was built — and how to check it

Jenkins's programme of enquiry

The observations in Sections 1, 2 and 4 are drawn from a sustained, multi-decade programme combining scholarly synthesis, diagnostic measurement and applied field research: a longitudinal, mixed-methods action-research design extending from diagnostic research begun in 1985, through a research-based Master's (Honours) thesis completed in 1998, to the collaborative work of the Mutuality — The Future of Trust Project. It interweaves three strands: a systematic survey of previously completed but disconnected research; iterative action research in which theoretical relationships were tested and applied in practice; and the development of a diagnostic instrument, the Essential Skills Profile, to measure applied cognitive capability rather than acquired knowledge. Data were generated in a dedicated learning centre and laboratory, where clients were tested, remediated and retested, and later through corporate workshops and cognitive coaching validated by retesting and 360-degree feedback. Insights were derived through an iterative, abductive process of pattern recognition across sources, and the literature was treated as primary data in its own right through thematic annotated bibliographies. The QF32 case draws on de Crespigny (2012) and the Australian Transport Safety Bureau's investigation (2013).

The Fifth Revolution evidence bank

The figures in Sections 3 to 9 are taken as published in *The White Whale Report 2026* (Byrne, Handley and Klaassen, 2026) and resolve to Frozen Package N1500 v1 of the Fifth Revolution evidence bank: 1,500 quantitative cases, each carrying five 0–5 scores (AI acceleration, HI transition capacity, societal intelligence, legitimacy, consequence intensity), seven metadata fields, an evidence tier (A=1.0, B=0.75, C=0.5), an outcome path and a written rationale with a source locator. The bank grew from 200 cases in 2025 to 1,500 in September 2026; every new case since N=1,300 entered with a page-level locator verified live. It is not a random sample and does not support population prevalence claims; it over-represents English-language reporting and regulated sectors; scores are interpretive ordinal codes; phases are coded at the most recent observable position. Every contrast in this paper is a descriptive comparison between coded groups drawn from the same case records, and no causal or commercial effect is claimed. The arithmetic reproduces from the released files; data hygiene survived three adversarial audits at N=1,300 with corrections disclosed; and two independent validators are re-coding a blinded stratified sample of 225 cases, classifying every contested-phase case as designed or default before its outcome is known, and double-coding the accountability and capability readings. Case identifiers in this paper (for example N1500-068) resolve to page-level source locators in the bank's register. Reviewers who wish to examine the frozen package, propose cases or challenge a score are invited to write to Building Mutuality.

Disclosure. Both authors have a commercial interest in the practices this paper describes — Jenkins through the Essential Skills Profile and cognitive development programmes, Byrne through Building Mutuality and the HI Practitioner accreditation. The evidence bank is maintained separately from Building Mutuality's commercial instruments, and no case in it was scored with reference to the accreditation. Research and analysis were conducted with AI assistance under the authors' direction; every editorial judgement, scoring anchor and disclosure is the authors' responsibility.

About the authors

Denis Jenkins is a cognitive researcher, educator and corporate transformation specialist with 28 years of cognitive investigation and practice. His research thesis, *Journey Into Learning* (University of New England, 1998), for which he was awarded a Master's degree with Honours, and the Essential Skills Profile he developed with Helga Rowe underpin a body of work on cognitive development, information management and mutuality. He has worked with cognitive profiles across education, aviation and the corporate sector, and is co-founder of the Mutuality — The Future of Trust Project and co-author of *Mutuality: The Future of Trust* (2024). He coined the terms Adaptive Shared Intelligence and Human Adaptive Intelligence, and his first white paper (2026) proposed the shift in the concept of intelligence on which this paper builds.

Matthew Byrne is the founder of Building Mutuality, lead author of the Fifth Revolution research programme and Program Director of the AGSM General Manager Program at UNSW Business School. He created the Fifth Revolution evidence bank and is lead researcher of the Alignment–Mobilisation Model on which *The White Whale Report 2026* rests. He is the author of *The 5th Revolution: Human Intelligence in the Age of AI* (2026) and of the white papers and articles cited here, co-author with Jenkins of *Mutuality: The Future of Trust* (2024), and co-designer of the HI Practitioner accreditation. His work is about one question: how people stay needed, and know it, in the future their organisations are building.

Cite as: Jenkins, D. and Byrne, M. (2026) *Understanding the Pillars of Mutuality and Trust That Lead to Human Adaptive Intelligence: A Proposed Shift in the Concept of Intelligence*. White Paper 2. Mutuality — The Future of Trust Project, September 2026. When quoting an evidence-bank figure, cite the exhibit and Frozen Package N1500 v1; when quoting a case, cite its identifier.

References

Acemoglu, D. and Johnson, S. (2023) *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs.

Albanese, A. (2026) Interview, *Insiders*, ABC, 20 September 2026; reported as "Albanese at odds with Trump over AI rules, warns humans must be in control", ABC News.

American Psychological Association (2023) *2023 Work in America Survey: Artificial intelligence, monitoring technology, and psychological well-being*. Washington, DC: APA (n = 2,515; fieldwork April 2023).

Amodei, D. (2026) "We must pace the frontier", essay, 12 September 2026; responses by S. Altman and E. Musk as reported by TechCrunch and the Washington Post, 12–14 September 2026.

Australian Government, Department of Industry, Science and Resources (2024) *Voluntary AI Safety Standard*. Canberra.

Australian Transport Safety Bureau (2013) *In-flight uncontained engine failure, Airbus A380-842, VH-OQA, overhead Batam Island, Indonesia, 4 November 2010*. Investigation AO-2010-089. Canberra: ATSB.

Bainbridge, L. (1983) "Ironies of automation", *Automatica*, 19(6), pp. 775–779.

Baumeister, R.F. and Leary, M.R. (1995) "The need to belong: desire for interpersonal attachments as a fundamental human motivation", *Psychological Bulletin*, 117(3), pp. 497–529.

Becker, J., Rush, N., Barnes, E. and Rein, D. (2025) *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*. METR; arXiv:2507.09089.

Brynjolfsson, E. (2022) "The Turing Trap: the promise and peril of human-like artificial intelligence", *Daedalus*, 151(2), pp. 272–287.

Brynjolfsson, E., Li, D. and Raymond, L. (2025) "Generative AI at work", *Quarterly Journal of Economics*, 140(2), pp. 889–942.

Byrne, M. (2025) *HI–AI: Aligning Human Intelligence with Artificial Intelligence Integration*. White paper. Sydney: Building Mutuality.

Byrne, M. (2026a) *Why AI Feels So Scary to Humans: From Prediction to Possibility*. White paper, February 2026. Sydney: Building Mutuality.

Byrne, M. (2026b) *How AI Amplifies Your Culture*. White paper, March 2026. Sydney: Building Mutuality.

Byrne, M. (2026c) *Building HI Capability: What Human Intelligence Must Become When Prediction Is Abundant*. White paper, March 2026. Sydney: Building Mutuality.

Byrne, M. (2026d) *Mattering as the Psychological Core of Human Intelligence Transition Capacity*. Fifth Revolution Research Programme, July 2026. Sydney: Building Mutuality.

Byrne, M. (2026e) *The Intelligence Between Us: A Response to Denis Jenkins on Human Adaptive Intelligence*. Sydney: Building Mutuality.

Byrne, M. (2026f) *The 5th Revolution: Human Intelligence in the Age of AI*. Sydney: Building Mutuality (forthcoming, November 2026).

Byrne, M. and Jenkins, D. (2024) *Mutuality: The Future of Trust — The Executive's Handbook, Enabling Mutuality to Cognitively Transform Organisations*. Victoria, BC: Tellwell Talent.

Byrne, M., Handley, D. and Klaassen, R. (2026) *The White Whale Report 2026: The State of the Human–AI Transition. Evidence from 1,500 cases*. Edition 1. Sydney: Building Mutuality, The 5th Revolution research programme. Data: Fifth Revolution evidence bank, Frozen Package N1500 v1.

California Assembly Bill 316 (2025), Civil Code s 1714.46, in force 1 January 2026.

Capra, F. (1996) *The Web of Life*. New York: Anchor Books.

Chun, J.S., De Cremer, D., Oh, E.-J. and Kim, Y. (2024) "What algorithmic evaluation fails to deliver: respectful treatment and individualized consideration", *Scientific Reports*, 14, 25996.

Commonwealth Bank of Australia (2026) "Commonwealth Bank's $90m plan for AI-ready workforce", newsroom release, 24 February 2026.

Davies, D. (2024) *The Unaccountability Machine: Why Big Systems Make Terrible Decisions — and How the World Lost Its Mind*. London: Profile Books.

de Crespigny, R. with Abernethy, M. (2012) *QF32*. Sydney: Pan Macmillan.

Dehaene, S. (2020) *How We Learn: The New Science of Education and the Brain*. London: Allen Lane.

Dell'Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R. (2023) "Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality", Harvard Business School Working Paper 24-013.

Doidge, N. (2007) *The Brain That Changes Itself*. New York: Viking.

Elish, M.C. (2019) "Moral crumple zones: cautionary tales in human–robot interaction", *Engaging Science, Technology, and Society*, 5, pp. 40–60.

European Union (2024) Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 4 (AI literacy) and Article 14 (human oversight).

Fernandes, D., Villa, S., Nicholls, S., Haavisto, O., Buschek, D., Schmidt, A., Kosch, T., Shen, C. and Welsch, R. (2026) "AI makes you smarter but none the wiser: the disconnect between performance and metacognition", *Computers in Human Behavior*, 175, 108779.

Feuerstein, R. (1980) *Instrumental Enrichment: An Intervention Program for Cognitive Modifiability*. Baltimore: University Park Press.

Flavell, J.H. (1979) “Metacognition and cognitive monitoring: a new area of cognitive–developmental inquiry”, *American Psychologist*, 34(10), pp. 906–911.

Flett, G.L. (2018) *The Psychology of Mattering: Understanding the Human Need to Be Significant*. London: Academic Press.

Fukuyama, F. (1995) *Trust: The Social Virtues and the Creation of Prosperity*. New York: Free Press.

Gillespie, N., Lockey, S. et al. (2025) *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*. Melbourne: University of Melbourne and KPMG (n = 48,000+, 47 countries).

Helmreich, R.L., Merritt, A.C. and Wilhelm, J.A. (1999) “The evolution of Crew Resource Management training in commercial aviation”, *International Journal of Aviation Psychology*, 9(1), pp. 19–32.

Henrich, J. (2015) *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton: Princeton University Press.

Hodson, R. (2001) *Dignity at Work*. Cambridge: Cambridge University Press.

Holland, J.H. (1995) *Hidden Order: How Adaptation Builds Complexity*. Reading, MA: Addison-Wesley.

Hrdy, S.B. (2009) *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Cambridge, MA: Harvard University Press.

Humlum, A. and Vestergaard, E. (2025) “The unequal adoption of ChatGPT exacerbates existing inequalities among workers”, *Proceedings of the National Academy of Sciences*, 122(1).

Hutchins, E. (1995) *Cognition in the Wild*. Cambridge, MA: MIT Press.

Jenkins, D. (1998) *Journey Into Learning*. Research thesis, Master’s (Honours), University of New England, Australia.

Jenkins, D. (2026) *Mutuality and Trust Are the Foundation of Human Adaptive Intelligence: A Proposed Shift in the Concept of Intelligence*. White Paper 1. Mutuality — The Future of Trust Project.

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and 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”, in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. New York: ACM.

Malone, T.W. (2018) *Superminds: The Surprising Power of People and Computers Thinking Together*. New York: Little, Brown.

Matthias, A. (2004) “The responsibility gap: ascribing responsibility for the actions of learning automata”, *Ethics and Information Technology*, 6(3), pp. 175–183.

McKinsey & Company (2026) *The State of Organizations 2026*. New York: McKinsey & Company.

Meadows, D.H. (2008) *Thinking in Systems: A Primer*. White River Junction, VT: Chelsea Green.

Mercurio, Z. (2025) *The Power of Mattering: How Leaders Can Create a Culture of Significance*. Boston: Harvard Business Review Press.

Moffatt v Air Canada, 2024 BCCRT 149 (British Columbia Civil Resolution Tribunal, 14 February 2024).

Mulgan, G. (2018) *Big Mind: How Collective Intelligence Can Change Our World*. Princeton: Princeton University Press.

Nadella, S. (2026) Interview with the *Wall Street Journal*, June 2026; as reported by Fox Business, “Satya Nadella warns AI companies must earn public trust on job impact”.

Norman, D.A. (1993) *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, MA: Addison-Wesley.

Noy, S. and Zhang, W. (2023) “Experimental evidence on the productivity effects of generative artificial intelligence”, *Science*, 381(6654), pp. 187–192.

Ostrom, E. (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Otis, N., Clarke, R., Delecourt, S., Holtz, D. and Koning, R. (2023) “The uneven impact of generative AI on entrepreneurial performance”, SSRN Working Paper 4671369.

Prilleltensky, I. (2020) *How People Matter: Why It Affects Health, Happiness, Love, Work, and Society*. Cambridge: Cambridge University Press.

Putnam, R.D. (2000) *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.

Rosenberg, M. and McCullough, B.C. (1981) “Mattering: inferred significance and mental health among adolescents”, *Research in Community and Mental Health*, 2, pp. 163–182.

Royal Commission into the Robodebt Scheme (2023) *Report*. Canberra: Commonwealth of Australia.

Schön, D.A. (1983) *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books.

Senge, P.M. (1990) *The Fifth Discipline: The Art and Practice of the Learning Organization*. New York: Doubleday.

Siemiatkowski, S. (2025) Remarks on Klarna’s reinvestment in human customer support, as reported by CX Dive (2025); evidence-bank case N1000-017.

Stanovich, K.E. (2016) *The Rationality Quotient: Toward a Test of Rational Thinking*. Cambridge, MA: MIT Press.

Sweller, J. (1988) “Cognitive load during problem solving: effects on learning”, *Cognitive Science*, 12(2), pp. 257–285.

Tomasello, M. (2019) *Becoming Human: A Theory of Ontogeny*. Cambridge, MA: Harvard University Press.

United States House of Representatives, Committee on Transportation and Infrastructure (2020) *Final Committee Report: The Design, Development and Certification of the Boeing 737 MAX*. Washington, DC, September 2020.

Vygotsky, L.S. (1978) *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press.

Wang, R.E., Ribeiro, A.T., Robinson, C.D., Loeb, S. and Demszky, D. (2024) *Tutor CoPilot: A Human–AI Approach for Scaling Real-Time Expertise*. Stanford University; arXiv:2410.03017.

Weick, K.E. and Sutcliffe, K.M. (2007) *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*, 2nd edn. San Francisco: Jossey-Bass.

Wood, D., Bruner, J.S. and Ross, G. (1976) "The role of tutoring in problem solving", *Journal of Child Psychology and Psychiatry*, 17(2), pp. 89–100.

Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N. and Malone, T.W. (2010) "Evidence for a collective intelligence factor in the performance of human groups", *Science*, 330(6004), pp. 686–688.