

From Chatbots to Agents

Leading your organisation into the AI-underpinned era. Why delegating work to AI is a different operating model from prompting a chat box, and how to make the transition without repeating history's mistakes.

Turn Intent Into Impact.



The chatbot era: powerful, but human-paced

The first wave of business AI lived in a chat box. You asked, it answered, you read, you asked again. Useful, but every output still needed a human to drive every step.

~80%

of ChatGPT conversations were guidance, information seeking and writing

NBER working paper 34255, with OpenAI data

2 of 3

work writing requests were edits to existing text, not new work

NBER working paper 34255

78%

of organisations used AI somewhere, mostly chat and content drafting

McKinsey State of AI, 2025

"The true value of AI lies not in whether the tool can perform a task. It almost always can. It lies in the quality of the idea driving it."

Four rungs from tool to colleague

Each rung hands the machine more autonomy. The chatbot era stopped at rung two. The question at every step: how much do you delegate, and how do you review?

1 Automation

Rules run the task. Same input, same output. No judgment involved.

2 Assistant

You drive, it helps. The whole chat-box era lives on this rung.

3 Agent

You delegate an outcome. It plans, uses tools, and returns finished work.

4 Agentic

Standing systems of agents run whole processes. Humans review at checkpoints.

THE PIVOT, MEASURED

Chat is collaboration. Deployment is delegation.

Anthropic's own usage data tells the story in one split. On the consumer chat product, automation and augmentation run roughly even. But when businesses embed AI into their operations, they stop chatting and start delegating whole tasks.



49 / 47

consumer chat: automation vs augmentation, roughly even

77%

of business API transcripts show full task delegation

12%

show back-and-forth collaboration

Source: Anthropic Economic Index, September 2025 report. Capability trendline: METR, task length agents complete autonomously doubles roughly every 7 months.

What agentic usage actually looks like



Delegate like a manager

Brief an agent the way you would brief a capable new starter. The intern test: could you teach and trust an intern to do it? Then you can delegate it. Prompt quality is really management quality.



Autonomy with checkpoints

Agents work in the background for hours and return finished work for review. Microsoft calls this the agent boss era: every worker manages agents, and leaders set the human-to-agent ratio per task.



Review is structural

Systems propose; people decide. Most production teams keep agents read-only or require human approval before anything is written, sent, or deleted. Review is part of the design, not a courtesy.




Proof of practice: the author rebuilt his own company as an autonomous enterprise in two months. Agents run content, outreach, meetings and operations on his own hardware. Parts of this very document were produced by agents working under human review.

Loop engineering

The discipline of designing the repeating loops your agents run in. A one-off delegation gets a task done. A standing loop gets a business process done every day without you.

Every loop has the same five beats, and each beat maps to a management skill you already have: write the brief like a product spec, let the agent work, review the 20% where human judgment adds the value, feed back what was wrong, and version the brief like a prototype.

"Systems propose; people decide."



Standing loops cannot be allowed to fail silently

WHICH PROCESSES BECOME LOOPS

- Recurring work you can describe as rules and check against evidence: reporting, triage, follow-ups, content drafts, monitoring
- Run Eliminate, Automate, Delegate first. The most impactful move is often elimination, not automation
- Keep one-off delegations one-off. Not everything that can be automated should be

Frameworks: Eliminate–Automate–Delegate, the intern test, the 80/20 Collaboration Rule.

LOOP HEALTH: THE NON-NEGOTIABLES

- Agents rarely crash outright. Failure arrives as degraded output and silent drift, so define what silent failure looks like and who checks
- Audit logs for every AI decision, reversibility for every automated action, default to human judgment in ambiguity
- Layer the watch like Three Lines of Defence: operators, then risk owners, then audit. 89% of organisations now run some agent observability

After the 2010 Flash Crash, markets got circuit breakers. Give your loops the same.



The cautionary tale: Klarna's AI handled 2.3 million chats, claimed the work of 700 agents, then the CEO publicly rehired humans because cost had been prioritised over quality. Delegation without review loops is not automation. It is abdication.

Factories swapped steam for electricity and gained nothing for decades

Paul David's famous dynamo research: factories first replaced their steam engine with a dynamo and changed nothing else. The productivity gains only arrived when they physically and organisationally redesigned the factory around the new power source. Bolting agents onto unchanged workflows is the same mistake, a century on.

Under 5%

of US factory power was electric twenty years after the technology was commercially viable

~40 years

from Edison's first power station to measurable factory productivity gains

70%

of digital transformations still fall short of their objectives today

Sources: Paul David, The Dynamo and the Computer, American Economic Review 1990. Brynjolfsson, Rock and Syverson, The Productivity J-Curve, 2021: general purpose technologies demand workflow redesign and retraining before measured gains appear. BCG, Flipping the Odds of Digital Transformation Success, 2020.

A cultural transplant, not a tech update

Skills

The Five Cs become team protocols: Communication, Creativity, Critical Thinking, Collaboration, Curiosity. Teams learn to blend carbon and silicon colleagues.

Management

From doing the work to orchestrating it. Performance shifts toward how well people brief, review and improve their agents. Moderna merged HR and IT under one leader to plan people and agents together.

Trust

Only around 40% of workers have had any AI training, most rely on AI output without evaluating it, and 61% hide their AI use. Shadow AI thrives wherever leadership stalls.

Accountability

A computer can never be held accountable, so somebody must be. Fairness, Accountability, Transparency baked in, with humans owning every decision that matters.

THE UNCOMFORTABLE FINDING

Employees are ready. Leadership is the bottleneck.

McKinsey surveyed thousands of employees and hundreds of executives. Only 1% of leaders call their company mature in its AI use. Meanwhile employees already use AI at three times the rate their leaders estimate. The workforce is not waiting for permission. The organisation chart is.

"The people best at AI are not the most technical. The people who struggle most are the ones waiting to be told the exact way to do something."

Sources: McKinsey, Superagency in the Workplace, January 2025. Quote: BBE J-Wiki, observation from delivery.

10%

A practical path for an SME

DAYS 1-30

Map and set the rails

Run Eliminate, Automate, Delegate over your processes. Use the intern test to shortlist three delegation candidates and write each brief like a product spec. Before any agent runs: audit log, reversibility, human sign-off.

DAYS 31-60

One standing loop, piloted

Launch one loop in sandbox conditions: structured, safe, reversible. Define its health signals, agree who checks them, and apply the 80/20 review to every output.

DAYS 61-90

Prove it, then scale it

Measure outcomes, not activity. Decide Buy versus Build for loops two and three using the Four Bs: Bandwidth, Bootstrap, Buy, Build.

57% to 86%

daily AI usage on one client programme

9.7 / 10

programme NPS

133 hrs

returned per month by 24 automated processes

Keep yourself honest: Gartner predicts over 40% of agentic AI projects will be cancelled by the end of 2027 for unclear value and weak risk controls. The 90-day discipline above exists to keep you out of that statistic.

Three questions for your next board meeting

01 When AI decides, who carries the liability?

02 Are we supporting the human transition, or just announcing a system change?

03 If expertise is free, where is our edge?

"Machines machine better than people. People need to people better than machines."

Built on evidence and lived practice

Anthropic	Economic Index, September 2025 report	anthropic.com/research
NBER	How People Use ChatGPT, working paper 34255, 2025	nber.org/papers/w34255
Microsoft	2025 Work Trend Index: The Frontier Firm	microsoft.com/worklab
McKinsey	Superagency in the Workplace, January 2025; State of AI, November 2025	mckinsey.com
Paul David	The Dynamo and the Computer, American Economic Review, 1990	aeaweb.org
Brynjolfsson et al.	The Productivity J-Curve, AEJ: Macroeconomics, 2021	aeaweb.org
KPMG / Univ. of Melbourne	Trust, Attitudes and Use of AI: A Global Study, 2025	kpmg.com
Gartner	Agentic AI project cancellation prediction, June 2025	gartner.com/newsroom
METR	Measuring AI Ability to Complete Long Tasks, March 2025	metr.org
LangChain	State of Agent Engineering, 2025	langchain.com
Press coverage	Klarna reversal (Entrepreneur, CX Dive 2025); Moderna HR-IT merger (CIO, WSJ 2025)	various
Bykov-Brett Enterprises	Training sessions S1 to S10 and the nine BBE frameworks, J-Wiki	bykovbrett.net

Lead the transition. Don't just announce it.

We design and run agentic operating models for SMEs, on evidence, not hype. If you want help mapping your first standing loops, we have done it for our own business first.

Talk to us: bykovbrett.net/contact

More guides: bykovbrett.net/resources



Turn Intent Into Impact.