Work in progress. Perspective, 6 November 2026. Not peer-reviewed research.

The enterprise is where AI co-existence gets real

Mr. Gómez · Perspective · 6 November 2026

Abstract

A work-in-progress perspective, not a peer-reviewed paper, not an empirical study, and not a completed account of displacement. It responds to the published excerpt and publisher descriptions of Ethan Mollick's Co-Existence, not to a reading of the full book. I treat large enterprises as an early setting in which authority, work and accountability have to be rearranged together, not as the first users of AI and not as a representative sample. The argument is that organisational fit is the missing piece: a faster calculation and a faster organisation are not the same thing. Insurance and retail scenes are anonymised conversations, not audited case studies. Published evidence is cited inside its original limits. The 2027 expectation is a bet that publication-cleared examples of bounded production autonomy will appear, and it is explicitly open to being wrong. Disclosure: I work at Strategy, formerly MicroStrategy. My interest in semantic definitions overlaps with my employer's business. The argument is mine, not a statement on its behalf.

WORK IN PROGRESS | Perspective, 6 November 2026

This is a working argument, not a research paper or a completed account of displacement. It will harden as publication-cleared process evidence arrives. I separate what I have heard, what published evidence supports and what I expect next.

Up-front disclosure: I work at Strategy, formerly MicroStrategy. My interest in semantic definitions overlaps with my employer's business. The argument here is mine, not a statement on its behalf.

A price table takes six to twelve months to change. The question is whether AI can make it move faster.

That is an insurance conversation I have encountered, not an industry benchmark. The time includes boards, committees and people, not just someone crunching numbers. I do not yet have a publishable breakdown of where those months go.[4]

That gap matters. A faster calculation and a faster organisation are not the same thing.

Co-existence has an organisational address

Ethan Mollick's *Co-Existence: The Next Phase of AI* gives this question a useful starting point. It was published on 20 October 2026 by Portfolio in the US and WH Allen in Australia.[1,2]

In the published Wharton excerpt, Mollick argues that current AI already has enough unused capability to drive significant change. Stronger systems may arrive, but his argument does not depend entirely on them. He also distinguishes technological change from the choices people make about deployment.[3]

I am responding to that excerpt and the publisher descriptions, not pretending this is a review of the full book.

My extension is organisational. I think big enterprises are society's first wave of AI co-existence at an institutional scale: places where we can study the effects as they unfold, rather than only predict them.

Not the first people to use AI. Not a representative sample of society. An early setting in which authority, work and accountability have to be rearranged together.

The Australian Bureau of Statistics offers a reason to look there. Around 35% of large businesses reported using AI in 2024–25, compared with around 12% of businesses overall.[5] Those figures describe AI use, not autonomous agents or removed jobs. They support the choice of setting, not the conclusion.

I want to know what happens between buying capability and changing an operating model.

The missing puzzle piece is fit

In conversations with retailers, banks and insurers, I keep coming back to the same problem. You have the AI piece, but do not know where it fits in the enterprise puzzle. Or you know which piece is missing, but cannot make the technology fit it.[4]

The frontier conversation about singularity and superintelligence does not describe the work I am hearing about. That is my reading of those conversations, not a claim that model capability has stopped mattering.

The practical questions are about change management. If this step becomes automated, who is affected? Which approval still matters? Does the next team accept the output? Do we still need the same headcount?

A model can be clever and the process can remain unchanged. My concern is not model smartness in isolation, but whether an organisation can absorb the capability without losing control of the work.

Published workplace evidence is useful precisely because it is narrower than a grand prediction. Brynjolfsson, Li and Raymond's NBER working paper reports a 14% average increase in issues resolved per hour among 5,179 customer-support agents given access to an AI assistant, with gains differing across workers.[6]

That is evidence about assistance in a particular workplace. It does not tell us that an autonomous replacement would work, or that every organisation would turn saved time into fewer roles.

The puzzle includes the task, the people and the permission to change both.

The insurance question is about elapsed time

In the insurance scene, the ambition is to change pricing faster. AI might accelerate the number-crunching behind a proposed table. Whether that changes a six-to-twelve-month cycle is still an open question.[4]

If the delay is analytical, better tools could help. If it is mostly waiting for a committee, producing the answer sooner may just put it in the queue earlier. If a committee is protecting a necessary decision right, removing it is not a technical optimisation.

I would not call all those months waste.

APRA's general-insurance guidance discusses pricing delegations, actuarial involvement, monitoring and responding to emerging experience. It does not establish a six-to-twelve-month repricing norm or require a board to approve every price table.[7]

That distinction stops my observation becoming a false industry rule.

What would persuade me? A cleared account showing the old elapsed time, the analytical time within it and the approvals that remained necessary. Then the same account after the workflow changed, including mistakes and rework.

Not just a faster model response, but a decision implemented sooner without weakening the controls that justified the process in the first place.

Until that evidence arrives, I can describe the question people are asking. I cannot claim they have answered it.

The retail question is about exceptions

The retail conversation is different. The ambition I have heard is an army of bots watching logistics and supply-chain problems, spotting delays and dealing with exceptions. Human hand-offs are part of the lag they want to remove.[4]

That is a proposed operating model, not a production result I can report.

To make the choice concrete, imagine a hypothetical exception workflow. A shipment is delayed. One agent checks the consequence; another prepares a response. The important boundary is whether either can change an order or commit money, rather than merely recommend what a person should do.

Those are illustrative steps, not details from a client implementation.

A human approval at every hand-off could preserve the original queue. Removing every approval could let a mistaken assumption travel further before anyone challenges it.

I would separate routine, reversible decisions from decisions with material consequences. I would also ask who gets called when several exceptions arrive together, and whether that person can actually intervene.

The point is not that humans always create delay and bots always remove it. It is that the human role needs a purpose other than standing beside each automated step.

An army of bots is still a management problem. Who sets the boundaries, notices a shared mistake and has the authority to stop the work?

Profit, productivity and the pull towards autonomy

My reading of the commercial pressure is simple: profit leads to productivity, which leads to questions about agent autonomy.[4] That is the incentive chain I am worried about, not a universal law of management.

If an agent does the work and a person repeats all of it, have we improved productivity or created double spending?

There may be good reasons to pay for that duplication during a transition. I do not object to it as a temporary control. I object to counting the machine's output as a saving while leaving the checking

effort outside the calculation.

The alternative being discussed is one human controlling many bots, perhaps ten, instead of one human beside each agent.[4] Ten is an illustration from those conversations, not an evidenced supervision ratio.

I would want to count escalations, correction, monitoring and recovery before calling that arrangement cheaper. An overseer with too little time or authority is not a meaningful safeguard.

For workers, the difference matters. Faster work might mean more output, shorter queues, redeployment or fewer positions. A task-level gain does not decide which outcome management chooses.

Jobs and Skills Australia's exposure analysis is an important counterweight. It finds augmentation potential generally outweighs automation with current technologies and stresses that employment effects depend on adoption and workplace adaptation.[8]

That challenges any easy story in which productivity must end in replacement. It does not remove the need to ask who receives the gains.

Co-existence is not a promise of permanent partnership

My own view is that we face a co-existence period while more computer-doable work becomes technically and organisationally replaceable. I do not know its length, or whether some roles will remain mixed indefinitely.

I would put work done entirely through software under particular scrutiny. But being able to reproduce its visible steps is not the same as being able to carry its obligations, handle its exceptions or justify replacing the person doing it.

I am not borrowing a dramatic quotation to settle that distinction. It is my argument, and it needs evidence.

Trust needs a definition before it needs a promise

In my conversations, much of the work is still copilot use: asking questions and getting information, while people remain unsure they can trust the answer.[4]

My concern is certified meaning. Before delegating a decision, I want an agreed definition of the thing the agent is deciding about. In a hypothetical retail workflow, does available stock include goods already committed elsewhere? That is a business definition, not a test of eloquence.

Strategy's Strategy Intelligence Platform (SIP) describes a semantic layer for common definitions and a knowledge graph representing a business ontology.[9] I mention it because it explains my perspective, with the employment disclosure above. A product description is not evidence that the product will deliver a particular outcome.

The vendor-neutral question is whether definitions are governed, maintained and usable by the systems acting on them. Certification should mean an accountable owner has approved the meaning, not that a polished answer has earned our confidence.

There is bounded evidence worth considering. A data.world-authored technical report tested 43 insurance-domain questions. It reports average overall execution accuracy of 16.7% through direct SQL generation and 54.2% through an ontology-mapped knowledge graph using SPARQL.[10]

That is a 37.5 percentage-point difference in that benchmark. It used generated sample data and changed the representation and query path. It is not independent proof that every semantic-layer product improves accuracy by that amount, and it did not measure monetary savings.

Gartner also forecasts accuracy and cost benefits from prioritising semantics by 2027.[11] I treat those as forecasts, not measured returns. The case still needs to be made in the workflow where the definitions will be used.

Better meaning does not itself authorise an action. I want both questions answered: is this the right interpretation, and is the system allowed to act on it?

Australia has choices, not a single timetable

My Australian concern is a tension between caution and competitiveness. We may move towards something closer to a European-style regulatory approach. I worry about unnecessary friction, but that is a possible direction and a concern, not a claim that Australia has adopted a particular regime.

I would judge a control by the harm it addresses and the work it requires. Some friction may be necessary. Some may be a poor substitute for a clear decision right.

Model choice belongs in the same conversation. My preference is to test Chinese-developed and other open-weight options for cheaper tiers, keeping expensive frontier capability where the task justifies it. That is a workload policy to evaluate, not a claim that nationality or openness guarantees lower cost or acceptable risk.

RouteLLM provides a limited research precedent. Its authors report substantial cost reductions while retaining a chosen performance level on benchmarks, with results varying by benchmark and weaker out-of-distribution routing without additional training data.[12]

That supports testing a routing policy. It does not give an enterprise a guaranteed saving after validation, retries, infrastructure and governance.

Local announcements need similar care. OpenAI announced its Australian initiative in December 2025, including training plans and an infrastructure memorandum. Anthropic announced the official opening of its Sydney office in April 2026.[13,14]

Those are different milestones. Neither demonstrates that an Australian process now runs autonomously or that a job has disappeared. I would not use an office opening as a proxy for operational maturity.

My 2027 expectation is a bet

In the work I am describing, 2026 feels like foundations, copilots and training people to implement agents. I have not yet seen the concrete production-displacement scene that would make my broader argument an observed result.[4]

My feel is that 2027 is when production agents start actively pushing humans aside in some of these workflows. It is an expectation, not a finding about Australia or a date on which computer-based work ends.

By the end of 2027, I would expect to see publication-cleared examples in the enterprise settings I am discussing where agents own bounded production steps previously performed by people, with less human effort per completed outcome after supervision and rework are counted.

If the evidence still amounts to copilots, demonstrations and unchanged approval queues, my timing was wrong. If roles remain but their tasks change, I should describe task displacement or redesign, not imply redundancies.

There is a serious counter-case. Gartner forecasts that over 40% of agentic AI projects will be cancelled by the end of 2027 because of escalating costs, unclear value or inadequate risk controls. Its same release forecasts growing autonomous decisions by 2028.[15]

Neither prediction is an Australian employment estimate. Together they warn me against treating production as the inevitable next step after a promising pilot.

The year is useful only if I let it expose the argument to being wrong.

What would change this view

I would soften the autonomy argument if sustained deployments showed better service and lower total cost while keeping substantial human involvement. That would suggest the partnership is productive, not merely unfinished automation.

I would delay my 2027 expectation if integration, permissions and accountability remain unresolved despite adequate task capability. That would strengthen the fit argument while weakening my proposed timetable.

I would challenge the enterprise-first lens if smaller organisations or public institutions produced clearer evidence of changed authority and work. Large enterprises are where I am looking, not where change must begin.

I would reject a claimed displacement result if it rested only on a restructuring announcement. I would want to know which work moved, what human effort remained and whether the organisation had simply renamed a cost-cutting decision.

And I would take worker experience seriously alongside throughput. Can people question an output? Do they retain the knowledge needed to recover when automation fails? Are they learning a different job, or merely carrying responsibility for decisions they cannot control?

Those are the questions I want process evidence to sharpen. Not another declaration that the future has arrived, but an account of what changed and who now carries the consequences.

What I'd tell another practitioner

1. Find the missing puzzle piece before upgrading the model. Name the decision, the queue and the person authorised to change the process. Faster analysis is not automatically a faster organisation.

2. Count the human work left around the agent. Include supervision, escalation and recovery. One person overseeing many bots is a proposition to test, not a saving to book in advance.

3. Put a condition on your prediction. Mine is 2027, with evidence of bounded production autonomy and reduced total human effort. If that evidence does not arrive, I need to change my view.

Disclosure and references

I work at Strategy, formerly MicroStrategy. The product mention above is limited to its public description, not a comparative endorsement. The insurance and retail scenes are anonymised

accounts of conversations, not audited case studies. No confidential implementation details or unpublished experiment results are included.

1. Penguin Books Australia. *Co-Existence: The Next Phase of AI* (https://www.penguin.com.au/books/co-existence-9780753562369). Ethan Mollick; WH Allen; 20 October 2026.

2. Penguin Random House. *Co-Existence* (https://www.penguinrandomhouse.com/books/813026/co-existence-by-ethan-mollick/). Portfolio; 20 October 2026.

3. Mollick, E. The next phase of AI (https://knowledge.wharton.upenn.edu/article/the-next-phase-of-ai/). *Knowledge at Wharton*, 20 October 2026. Published excerpt, not the full book.

4. Gómez, C. Author's anonymised conversations, 6 November 2026. Unpublished first-person observations, not independently verified and not a public corroborating source.

5. Australian Bureau of Statistics. Business adoption of Artificial Intelligence accelerates in 2024–25 (https://www.abs.gov.au/media-centre/media-releases/business-adoption-artificial-intelligence-accelerates-2024-25). 25 June 2026.

6. Brynjolfsson, E., Li, D. and Raymond, L. R. Generative AI at work (https://www.nber.org/papers/w31161). NBER Working Paper 31161, 2023. Figures here refer to the working-paper version, not the later journal version.

7. APRA. GPG 240 Insurance Risk (https://www.apra.gov.au/practice-guides/gpg-240). Current 4 December 2024; pricing guidance.

8. Jobs and Skills Australia. Our Gen AI transition: exposure (https://www.jobsandskills.gov.au/studies/generative-artificial-intelligence-capacity-study/our-gen-ai-transition-exposure). Potential task effects, not realised job outcomes.

9. Strategy. Strategy Intelligence Platform (https://software.strategy.com/ssis/platform). Vendor description of semantic-layer and ontology services, accessed 6 November 2026.

10. Sequeda, J. F., Allemang, D. and Jacob, B. A benchmark to understand the role of knowledge graphs on large language models' accuracy for question answering on enterprise SQL databases (https://arxiv.org/abs/2311.07509). Vendor-authored technical report, 13 November 2023.

11. Gartner. Lack of semantics causes inaccurate AI agents and wasted spending (https://www.gartner.com/en/newsroom/press-releases/2026-05-11-gartner-says-lack-of-semantics-causes-inaccurate-artificial-intelligence-agents-and-wasted-spending). 11 May 2026; analyst forecasts, not measured savings.

12. Ong, I. et al. RouteLLM: an open-source framework for cost-effective LLM routing (https://www.lmsys.org/blog/2024-07-01-routellm/). Authors' research summary, 1 July 2024.

13. OpenAI. Introducing OpenAI for Australia (https://openai.com/global-affairs/openai-for-australia/). 4 December 2025; announcement and plans, not delivery verification.

14. Anthropic. Sydney office opening and ANZ general manager appointment (https://www.anthropic.com/news/theo-hourmouzis-general-manager-australia-new-zealand). 27 April 2026; company announcement.

15. Gartner. Over 40% of agentic AI projects will be canceled by end of 2027 (https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-

percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027). 25 June 2025; global forecasts.

~ Mr. Gómez