



Agentic AI Is Restructuring Work. Don't Be the Un-electrified Factory.

Autonomy Platform · Public Blog

PERSPECTIVE

© 2026 Azirella Ltd. All rights reserved worldwide.

[Read more at azirella.com/blog](https://azirella.com/blog)

Agentic AI Is Restructuring Work. Don't Be the Un-electrified Factory.

Every general-purpose technology since the printing press has restructured work. The firms that accepted the restructuring early captured the productivity gain. The firms that adopted the new technology in the form factor of the old one became the case studies the next generation reads. Agentic AI is running the same play right now.

The action this post is asking you to take:

Accept that agentic AI is a restructuring of work, not a faster way of doing the existing work.

New tools with old rules do not work. Old tools with new rules do not work.

The productivity surge belongs to the firms that change both, in the right order, the sooner the better.

The shape of the change is captured most precisely by Jensen Huang's task-versus-purpose distinction, which he has been making in nearly every recent appearance. In his Carnegie Mellon commencement address in May

2026, he gave it concrete form: a radiologist's task is reading scans; her purpose is caring for the patient. AI automates the task. It elevates the purpose. Tasks get automated; humans still own outcomes. The concrete operational expression of that elevation is what Huang said at NVIDIA's GTC announcement earlier this year:

“Employees will be supercharged by teams of frontier, specialized, and custom-built agents they deploy and manage.”

, Jensen Huang, NVIDIA · [GTC, March 2026](#).”

The role change is not the planner-with-an-AI-tool. It is the planner *as the manager of a team of agents*. The work shifts from doing the work to deploying, governing, and overriding the agents that do the work. And the economic stake of getting that transition right is the line Huang gave *Fortune* a month later:

“You won't lose your job to AI. You'll lose it to your coworker who uses it.”

, Jensen Huang, NVIDIA · [Fortune, April 2026](#).”

In a supply chain, the planner's work is moving from doing tasks to managing decisions. From processing 847 weekly exceptions to governing the agents that process them. From accumulating institutional knowledge in their head, to depositing judgement into a substrate that holds it. From firefighting on Monday morning, to inspecting the fourteen decisions where their judgement is most likely to improve the outcome. None of that role change happens, none of the productivity gain that follows it arrives, if leadership has not yet accepted that the work itself is being restructured.

The pattern is old enough to be useful. Every general-purpose technology that has reshaped industry and society since the printing press has had the same shape. The technology arrives. Firms that accept it as a restructuring of work, and redesign their processes and roles around what the new technology actually makes possible, capture the productivity surge. Firms that adopt the new technology in the form factor of the old one, trying to do the old work faster, do not. Electricity ran that experiment over thirty years. The personal computer ran it over fifteen. Agentic AI is running it now, and the lag is shrinking, but only at firms whose leaders are recognising the pattern early.

George Stalk's 1987 BCG paper *Time, the Next Source of Competitive Advantage* and the 1990 book *Competing Against Time* are the deeper argument for *why* the restructuring is worth doing. Stalk showed that a unit of physical product spends between 95 and 99.95 percent of its time in a corporate system waiting, and that firms compressing the wait compounded advantage on cost, quality, service, and capital intensity simultaneously. The argument transposes cleanly to decisions, which is the variable that defines competitive position in the next decade. The most direct treatment is in [The Decision Flow Problem](#). The platform that expresses it is at azirella.com. The path to the conversation we are built for sits at the foot of this page.

If you already agree with the action and want to skip the case, stop here. The rest of the post is the historical and operational record of the correct ordering for adopting a disruptive technology: **Understand the change** → **Adapt the process** → **Train the people** → **Select the technology**, and why no other ordering has ever delivered the productivity surge a new general-purpose technology promises.

The HBR *Insider* newsletter landed in my inbox this week with two pieces sitting next to each other that, taken together, describe what the next eighteen months of enterprise AI adoption will actually feel like.

The first piece, by BCG directors Julia Dhar, Kristy R. Ellmer, and Philip Jameson, is about a behavioural trap they call *false alignment*. Executive teams launch a transformation without explicitly agreeing on why, on what is and is not changing, and on how the changes will occur. They mistake “we have discussed this at least once” for agreement. Paralysis, hyperactivity, or tunnel vision follow. Failure becomes the most likely outcome.

The second piece, by Patrick van Esch, Yuanyuan Gina Cui, and J. Stewart Black, names a more specific failure mode. They call it the *agentic convergence trap*. It occurs when your AI and your competitors’ AI are deployed in the same market on the same platform, learning from each other, and over months and quarters they converge to the same decisions. You did not intend to cede strategic variation. You ceded it by adopting an AI platform under time pressure, accepting the default settings because questioning them required expertise the leadership team did not have, and quietly removing the human review process that had previously introduced variation, because removing that friction was the point.

The authors are explicit: this is not a technology problem. It is a leadership and governance one.

Both pieces are saying the same thing in different vocabulary. The hard part of an AI transition is not the AI.

A conversation with a former Gartner analyst

A few weeks ago I had a long conversation with someone I have known for years, a former senior analyst at Gartner. Over many years the analyst saw enterprises adopt and discard every major wave of supply-chain software since the early ERPs. The pattern is muscle memory by now.

The analyst's ordering of People, Process, and Technology was *Process. Then People. Then Technology.*

The reasoning is precise. You redesign the process first, because until you know which decisions get made, by whom, when, and based on what, you do not know what you are automating. You think about the people second, because once the process is redesigned you can identify which skills the people running it need. You select the technology last, because the technology only makes sense once you know what process it is supposed to support.

It is the order most consultancies wish they had the discipline to follow and most procurement departments wish they had the patience to allow.

With one important caveat, it is also the order that history has been quietly insisting on for five hundred years.

The caveat: steady state versus disruptive

Process. Then People. Then Technology. is the right ordering when the technology is a known quantity. You can write the requirements, select the implementation, slot the new tool into the redesigned process you already know how to design. That is the steady-state case, and it is the case most

enterprise technology decisions actually are. SAP to SAP. APS upgrade. ERP migration. The technology's shape is well-understood; what changes is the process you wrap around it.

A disruptive technology breaks this ordering, because the technology itself is what redefines the process. You cannot redesign the planner's role until you understand that agents can act on every decision within guardrails, twenty-four hours a day, with a near-zero marginal cost per decision. You cannot compress decision flow until you have absorbed that compression is now mechanically possible at a volume no team of human planners could ever have produced. You cannot move the planner from *doer* to *manager of agents* until you understand what an agent actually is, what managing one looks like, and what the override workflow does and does not capture.

So with a disruptive technology, an *Understanding* step gets inserted at the front of the sequence, and each subsequent step takes on a specific verb that makes its job explicit. The full ordering becomes:

Understand (the change) → Adapt (the process) → Train (the people) → Select (the technology).

Understand the change. Absorb what the technology now makes possible that wasn't possible before, and what it now makes impossible that used to define how the work was done. This is not technology selection; it is technology comprehension. The leadership team has to do this work itself; it cannot be delegated to procurement.

Adapt the process. Redesign the work around the new possibility, not around an idealised version of the old work, faster. The autonomous tram running on the same rails is the process that wasn't adapted. The factory

floor reorganised around individual electric drives at each machine is the process that was.

Train the people. Identify the new skills the adapted process actually requires. Override quality, governance of the agents, strategic variation against the convergence trap, the work that humans do when the agents are doing the doing.


Select the technology. Procurement, configuration, deployment. The step every steady-state playbook starts with, and the only step that should come last when the technology is disruptive. By the time you are here, you know what the process needs and the technology is a question of which implementation supports it.

The most common failure mode of disruptive-technology adoption is skipping the Understanding step entirely and trying to run a steady-state Process-People-Technology playbook on a technology whose implications the team has not yet absorbed. You get the autonomous tram on the same rails. The electric motor driving the same overhead shaft. The PC on the desk running the same paper-form-replicated workflow. The historical record is full of the firms that did this, and the firms that didn't.


To make the failure mode legible without a hundred years of hindsight, the cleanest analogy is one a friend offered me recently. Without understanding what the technology enables, you may decide to cross the Atlantic in a row boat. The steady-state ordering applied to that decision is reassuring. Choose the right row boat. Select rowers with endurance and muscle. Adapt the cadence to the wave conditions. All three are necessary, and all three are entirely beside the point. The question was never which

row boat to choose. It was whether the row boat is the right technology at all. The Understanding step is what surfaces that question. The other three steps answer it.

It is the order that history has been quietly insisting on for five hundred years.



The same point can be put on a different axis, which is the one most operators find more memorable.



New tools with old rules do not work. Old tools with new rules do not work. The diagonal is the whole argument.


The lag, every time

Every general-purpose technology that has reshaped industry and society has had a productivity lag. The technology arrived. The productivity gains did not. Years passed. Eventually some firms redesigned around the new technology, and at those firms productivity surged. The rest waited a generation, or went out of business.

The shape of the curve is, near enough, the same every time.

THE SHAPE OF EVERY TECHNOLOGY TRANSITION

Productivity over time. Most firms live on the plateau. The redesign is what unlocks the surge.



Time-aligned to the moment the technology becomes economically available. The horizontal distance between Phase 1 and Phase 3 is the lag. The vertical distance is what the redesigning firm captures and the non-redesigning firm forfeits.

A few specific cases, drawn from the standard economic-history literature.

 Jost Amman's 1568 woodcut of a printing house, showing a compositor at the type case and a pressman pulling the lever of a hand press.

Jost Amman, Buchdrucker (Printer), from Das Ständebuch, 1568. Public domain by age. Source: [Wikimedia Commons](#).

The printing press. Gutenberg's press is usually dated to around 1440.

The Protestant Reformation began in 1517, seventy-seven years later. The scientific revolution, with its journals and academies, took even longer. For the first fifty years the press mostly printed Bibles and indulgences in the form factor of the previous technology, the hand-copied manuscript. It took new institutions, the Lutheran congregation, the chartered academy, the public school, before the press's productivity arrived. The press was necessary. It was not sufficient.

Steam and the factory system. James Watt's improved engine was patented in 1769. The productivity surge associated with the factory system, the division of labour, urban migration, and the new managerial roles that ran the factories did not become the modal manufacturing form until well into the nineteenth century. Steam, like the press, needed an organisational form to express its productivity.

Railroads. Alfred Chandler's *The Visible Hand* (1977) is the canonical argument that the railroad created the modern professional management hierarchy. Standardised time zones, telegraph-based dispatch, interchangeable parts, the divisional structure with a head office, the M-form corporation - all of these were inventions of process and organisation that the railroad's technical capability demanded but did not directly provide. The railway took roughly forty years from the Liverpool & Manchester (1830) to mature managerial practice.

Electricity. Paul David's 1990 paper *The Dynamo and the Computer* documents the case in detail. By the 1890s, electric motors were widely available, and factories began installing them. They installed them as drop-in replacements for steam-driven central shafts and belt systems, with a single large electric motor driving the same overhead line shaft. Productivity barely moved. Two decades later, the productivity surge arrived, and it came from a redesign: factories were rebuilt around individual electric drives at each machine, with the factory floor reorganised by workflow rather than by proximity to the central shaft. The productivity gain was in the redesign of the factory, not in the electrification of the same factory. From the availability of electric motors to the productivity surge took roughly thirty years.



Interior of a late-nineteenth-century factory floor. Workers stand at individual machines beneath a forest of leather belts hanging from a single overhead line shaft that runs the length of the room. The shaft is driven by a central engine outside the frame.

A belt-driven factory floor, late nineteenth century. Every machine in the room takes its power from a single overhead shaft. When electricity arrived, factories like this one installed a single large electric motor in place of the steam engine and changed nothing else. The productivity surge came thirty years later, when the shaft was retired and each machine got its own electric drive. Public domain by age. Source: Library of Congress.

The telephone. Patented in 1876. For the first decade it was sold to businesses as a faster telegram, with operator-mediated calls between offices on the same urban exchange. The reorganisation of business around the telephone, customer service as a function, the centralised switchboard, the field sales force that could check in with headquarters mid-trip, was a slow process across the late nineteenth and early twentieth century.

The automobile. Ford's Model T arrived in 1908. The mass-car society, with its highways, suburbs, drive-through restaurants, drivers' licences, gas-station networks, and the entire insurance industry that supports it, was not a 1908 phenomenon. It was a 1950s and 1960s phenomenon. The technology was forty years ahead of the social and process redesign that absorbed it.

Powered flight. The Wright brothers in 1903. The commercial aviation we recognise, with air-traffic control, hub-and-spoke routing, scheduled airlines, the FAA, and aviation insurance, took fifty years to assemble. The aircraft was the easy part.



John T. Daniels's photograph of the Wright Flyer's first successful flight, 17 December 1903, with Orville Wright at the controls and Wilbur running alongside.

Orville Wright in flight, Kitty Hawk, North Carolina, 17 December 1903. Photographer: John T. Daniels.

The first powered, controlled, sustained flight in a heavier-than-air machine, twelve seconds.

Commercial aviation as we know it, with air-traffic control, scheduled hub-and-spoke routes, the FAA, and the insurance industry that underpins it, took half a century to assemble. Public domain by age.


Source: [Library of Congress, Prints and Photographs Division](#).

The personal computer. Robert Solow's 1987 review in *The New York Times Book Review* contains the line that has been cited more than any other in this debate: *we see computers everywhere except in the productivity statistics*. The PC was a 1981 product. The productivity surge associated with information technology in the US economy did not arrive until the late 1990s. Erik Brynjolfsson's productivity-paradox research showed why. Firms that simply installed computers got nothing. Firms that redesigned business processes around what computers made possible, with the people retrained accordingly, captured the gain. The lag was approximately fifteen years.

The pattern is consistent enough to be useful as a forecasting tool. The technology arrives. It is initially used in the form factor of the previous technology. Productivity disappoints. Some firms redesign their processes and roles around what the new technology actually makes possible. Productivity surges, but only at those firms. The new form becomes the new normal a generation later. The firms that read the pattern correctly start with the process redesign and let the technology follow.

THE LAG IS NEVER IN THE TECHNOLOGY

From the arrival of a general-purpose technology to the productivity surge it eventually delivers.



*The technology is the trigger. The redesign of process and people is the cause. Sources: Paul David, *The Dynamo and the Computer* (1990); Alfred Chandler, *The Visible Hand* (1977); Robert Solow, *NYT Book Review* (1987); Erik Brynjolfsson, *Productivity Paradox series* (1993–2003).*

The most recent observation tightens the trend. The PC lag was about fifteen years, faster than electricity or the automobile. Internet adoption was faster still. Mobile faster again. The lag is shrinking with each successive general-purpose technology, because organisations have become slightly better at recognising the pattern. But the lag is not zero, and the firms that read the historical pattern correctly are the ones that gain the most.

Every generation has its un-electrified factory.

The pattern is older than electricity, and it has never spared the firms that ignored it.

Where agentic AI is, right now

By the historical analogy, agentic AI is currently in the *form-factor-of-the-previous-technology* phase. The dominant deployment shape is a chat panel bolted onto a planning UI, with a frontier LLM summarising the screen back to the user in fluent English. This is the agentic-AI equivalent of a 1900s factory with one large electric motor driving the same overhead shaft. It works. It does not deliver the productivity gain that the technology actually makes possible.

The agentic convergence trap from the HBR piece is the form of disappointment we should expect. Vendors and customers both deploy the same AI in the same shape. The AI learns its way to the same local optimum that every other AI is finding. Strategic variation disappears, not because anyone consciously gave it up, but because no one redesigned the process to demand it.

False alignment is the form of failure at the leadership level. The team agrees that agentic AI is important, then disagrees implicitly about whether the project is supposed to cut costs, increase service level, capture institutional knowledge, replace planners, free planners up, or some combination. Each interpretation produces a different process redesign. None of them get done.

The way through is the same way through every previous technology transition. Redesign the process. Then identify the people skills the process needs. Then select the technology.

The clearest present-day analogue, in transit rather than supply chain, is the autonomous tram. Siemens deployed one in Potsdam, Germany, in 2018, a Combino tram retrofitted with cameras, lidar, radar, and on-board

reasoning sufficient to detect pedestrians, read traffic signals, and execute braking and acceleration decisions in real time. By every technical measure the autonomous tram is genuinely autonomous. The AI on board works.

What changes about a transit network when its trams become autonomous? Almost nothing, yet. The trams run on the same rails. They follow the same fixed routes. They stop at the same stations. They run on the same printed schedule. They serve the same passenger-flow patterns. The autonomy is a property of the vehicle. The network it operates in remains a fixed-rail urban-transit system designed in the late nineteenth century for human drivers on fixed routes. New tools. Old rules. The productivity gain of autonomous trams will arrive when the *network design* changes around the autonomous vehicle, when dynamic routing becomes possible, when fixed schedules give way to demand-responsive operations, when the fixed track itself is no longer the binding constraint. Until that happens, the autonomous tram is a more interesting commute on a less interesting transit system.

The supply-chain analogue is direct. An agent driving a fixed periodic-batch planning process is the autonomous tram. The agent is genuinely autonomous, but the system around it, the weekly run cycle, the human-throttled exception queue, the locked approval gates, the Monday-morning ritual of triage, is unchanged. New tools, old rules. The productivity gain arrives when the process is redesigned around what the agent makes possible, not before.

What the process redesign actually looks like

For supply-chain planning, the process redesign is fairly specific, and it is the redesign that has to come first.

Today, a senior supply-chain planner spends her week reviewing exception queues, making judgement calls under time pressure, and propagating the consequences of those calls across functions through phone calls and spreadsheets. Her week is set up around the assumption that she is the throughput-limiting step. She is also the institutional knowledge holder, which means when she retires, twenty years of pattern recognition leaves with her.

The redesigned process puts the planner in a different relationship to the work. The agents take every decision within guardrails overnight. They land their decisions in a structured queue, ranked by how likely the planner's judgement is to add value to that decision. The planner arrives Monday morning, looks at the top of the queue, and spends her time on the decisions where her override matters most. Every override is captured with reasoning. Every outcome is measured against what the agent would have done. Over time, two things happen. The agents get better at the decisions where the planner consistently improves outcomes. And the planner gets better at recognising the shape of the decisions where her judgement matters most, because she sees them and only them.

The role transformation is not "the planner uses an AI tool." It is "the planner is the manager of the agents." The agents do the work. The planner overrides and directs. The planner's value is in the override quality.

That is the process redesign. It is not a feature on a roadmap. It is the precondition.

THE PLANNER'S WEEK, BEFORE AND AFTER THE REDESIGN

Same supply chain, same Monday. The redesign moves the planner from doer to manager of decisions.

BEFORE THE REDESIGN

847

exceptions queued Monday

By Thursday: ~400 reviewed.
The other ~447 roll into next week.
Three of them are urgent.
One causes a stockout.

*Planner is the throughput-limiting step.
Knowledge walks out at retirement.*



AFTER THE REDESIGN

14

surfaced for inspection

847 decisions handled by the agent overnight.
14 at the top of the queue,
ranked by urgency × inverse-confidence.
Planner spends the morning on those.

*Planner manages decisions, not exceptions.
Every override is captured judgement.*

The 14 are not "the ones the agent could not handle." They are the ones the agent ranked highest on urgency × inverse-confidence, so that the planner's inspection time goes where her judgement is most likely to improve the outcome.

And then the people

Three skills follow from the process redesign, none of which the current planning workforce was hired or trained for.

The first is *override quality*. The planner gets better at knowing when to override and why. The system tracks override outcomes against what the agent would have done, and per-planner, per-decision-class, a reputation accumulates. Some planners will be net positive on substitution decisions and net negative on transportation routing. The system learns that, and weights training signal accordingly. Planners learn it about themselves, and play to their strengths. This is a different career than the one most planners signed up for, and a more interesting one.


The second is *governance*. Setting and adjusting the guardrails. Reviewing which decision classes are ready to graduate from Decision Support (the agent recommends, the human decides) through Augmentation (the agent decides within tight guardrails, the human inspects actively) to Automation (the agent decides within expanded guardrails, the human overrides as needed). Deciding when measured override-effectiveness over a rolling window has earned a decision class the right to graduate. This is the role of the planning leader, not the individual planner.

The third skill is the one the agentic convergence trap from HBR turns on. *Strategic variation*. If the agents are running the routine optimisation, the human contribution is the divergent intent. The willingness to choose differently than the algorithm proposes when the business strategy diverges from the operational optimum. The override is the structural place in the system where strategic variation is registered, captured, and fed back into the substrate the agents learn from. Companies that train their planners to override well, and reward them for doing so when their judgement is right, do not converge with their competitors. Companies that treat overrides as friction to drive down do.

This is what van Esch, Cui, and Black mean when they say the convergence trap is a leadership and governance problem, not a technology problem. The technology will converge by default. The process and the people are what create variation.

THREE NEW SKILLS THE REDESIGN DEMANDS

Each skill is anchored to a specific architectural commitment. None of the three is optional.



The technology supports these three skills. It does not replace them, and it does not produce them by default.

Three skills, three architectural anchors. Override Quality lives in the override loop. Governance lives at AI·IO·ML thresholds and the Decision-Support / Augmentation / Automation graduation. Strategic Variation is the reason the override is captured as a first-class signal in the substrate, not as friction to drive down.

And only then the technology

The technology selection falls out of the process and the people, not the other way around.

A platform supporting the redesigned process needs to do specific things. It needs to take decisions in real time without queueing on human approval, so the agent acts on every decision within guardrails. It needs to rank surfaced decisions by how likely the human's judgement is to improve them, so the planner's inspection time is spent where it matters. It needs to capture overrides with reasoning, classify them, and route them into the right place in the learning loop. It needs to graduate decision classes through deployment stages governed by measured override-effectiveness.

It needs a single canonical representation of operational reality that all the agents reason against, so they can be coordinated rather than competing. It needs to be auditable, so a regulated buyer can treat the platform's outputs as evidence rather than as oracle pronouncements.

These are not features. They are the technology consequences of the process redesign and the role transformation. The platform we are building at Azirella is the technology consequence of taking this sequence seriously. The order matters: the [AI·IO·ML operating model](#) is the process redesign. The [four channels of Operating Knowledge capture](#) are the people-skills artefact, the place where the new override-quality, governance, and strategic-variation skills are converted into training signal. The [Decision Stream](#) and the [eleven specialised agents](#) are the technology. The technology was selected last.

That is also why we are unenthusiastic about deployments that try to start in the opposite order. A chat panel bolted onto an existing planning UI does not redesign the process. It re-labels it. The convergence trap follows.

The closing thought, and back to the action

Every general-purpose technology since the printing press has had the same shape. The technology arrives. Firms accept it as a restructuring of work, redesign their processes and roles around what the new technology actually makes possible, and capture the productivity gain. Firms that do not accept the restructuring, that try to use the new technology to do the old work faster, do not capture the gain. They wait a generation, or go out of business. The lag has shrunk with each round, from thirty years for

electricity to fifteen for the personal computer, because organisations are getting marginally better at recognising the pattern. Agentic AI is the next round of the same play.

The role transformation Huang describes, from doer to manager of decisions, is the human-side expression of the restructuring. The process redesign is the operational expression. The technology is the consequence, not the cause. Acceptance is the precondition that lets any of it happen on a useful timescale.

The action this post asked of you at the top stands, in four steps with four explicit verbs.

One: Understand the change - what agentic AI now makes possible that wasn't possible before. Agents acting continuously within guardrails. Decision flow compressed by orders of magnitude. The planner moving from doer to manager of agents.

Two: Adapt the process - redesign the work around the new possibility, not around an idealised version of the old work, faster.

Three: Train the people - for the override quality, governance, and strategic-variation work that the adapted process actually demands.

Four: Select the technology that supports the adapted process.

The sooner you do step one, the sooner steps two through four pay off, and the further ahead you are of every competitor still running a steady-state procurement playbook on a disruptive technology. New tools with old rules do not work. Old tools with new rules do not work. The firms that change both, in the right order, are the ones the historical record rewards. The firms that don't become the case studies the next generation reads.

The post most worth your time next is [The Decision Flow Problem](#), where the Stalk parallel is set out formally with the BCG numbers and the underlying economics. The platform that expresses the argument in product form is at [azirella.com](#), the [Waiting](#) section applies Stalk's 0.05-to-5 percent rule directly to decisions, the [Planner's Day, Transformed](#) walkthrough shows what the redesign looks like when it lands, and the *Latency Compounds Costs. Velocity Compounds Value.* thread that runs through the whole site is the working assumption behind every architectural decision in the product. The path to the conversation we are built for sits at the foot of this page.

Further reading.

- Stalk, G. Jr. (1987). *Time, the Next Source of Competitive Advantage*. Harvard Business Review, July–August 1987. The original article. The full-length treatment is in the 1990 book with Thomas Hout, *Competing Against Time: How Time-Based Competition is Reshaping Global Markets*. Free Press.
- BCG (1987). *Rules of Response*. Perspectives series. Stalk's articulation of the 0.05-to-5 percent rule, the 3/3 rule, the 1/4-2-20 rule, and the 3-times-2 rule of compounding advantage from time compression.
- BCG, *How to Avoid the False-Alignment Trap When Leading Change* (Dhar, Ellmer, Jameson). Referenced in HBR *The Insider*, 18 May 2026.
- van Esch, P., Cui, Y. G., & Black, J. S., *The Agentic Convergence Trap*. HBR, May 2026.

- David, P. A. (1990). *The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox*. American Economic Review, Papers and Proceedings.
- Chandler, A. D. (1977). *The Visible Hand: The Managerial Revolution in American Business*. Belknap / Harvard.
- Solow, R. M. (1987). *We'd Better Watch Out*. New York Times Book Review, 12 July 1987 (the source of the productivity-paradox quote).
- Brynjolfsson, E. (1993–2003). *The Productivity Paradox of Information Technology* and follow-ups.
- Sandino, T., *How Fast-Growing Companies Can Make Better Decisions*. HBR, on structured empowerment.

See Autonomy in action

Walk through how Autonomy models, executes, monitors, and governs supply chain decisions with autonomous AI agents.

[See It Live](#)