HHSS Consulting

FROM FORDISM TO AI

The settlement after the production revolution

By Marcell Kisistók and Dhruv Agarwal

Strategic research report
September 2026

Executive summary

EXECUTIVE TAKEAWAY

AI has made output cheap. The scarce asset is increasingly the human judgement that verifies it and turns corrections into proprietary learning.

In January 1914 Henry Ford more than doubled the wages of his workforce, without a strike, a union, or any market pressure obliging him to. The decision is usually told as philanthropy or as public relations. It was neither. Ford's new assembly line had collapsed the cost of a car and, in the same movement, created an instability the line could not absorb: annual turnover of 370 percent and daily absenteeism of 10 percent in a production system where one empty station stopped everything.[1] The five-dollar day was a calculated purchase of the one thing his technology could not supply for itself.

This report argues that the AI economy of 2026 sits at the same point in the same sequence, and that the sequence is instructive rather than merely evocative. Generative AI has collapsed the cost of cognitive output, standardised the work that produces it, compressed the returns to experience, and generated its own characteristic instability. What it has not yet produced is a settlement: a deliberate management response aimed at whatever has become scarce.

Four findings carry the argument.

First, the binding constraint has moved. On the craft floor output depended on individual effort; on the assembly line it depended on retention; in AI-mediated work, where output volume is abundant and cheap, it depends on the human judgement that decides whether a given output is any good. Each production revolution changes what management is actually buying, and each time the change is recognised late.

Second, cheap AI is not an advantage. Ford's cost decline was proprietary and bought him roughly a decade of protected position. The AI cost decline is supplied by vendors and reaches every competitor simultaneously, and the performance gap between paid and freely available models has narrowed to 1.7 percent on benchmark tasks.[2] Falling costs pass through to prices and leave relative position unchanged. Advantage has to be built from what a vendor cannot supply.

Third, verification is that thing. The proprietary workflows, data, and feedback loops that constitute a durable moat are produced by people correcting model output, case by case, in a form the firm retains. Verification is therefore not an overhead on AI adoption; it is the mechanism by which adoption becomes advantage at all.

Fourth, the instability of this era is invisible in a way Ford's was not. He could count empty stations every morning. Cognitive deskilling, over-reliance, and unchecked output arrive as work delivered on time, and a firm that cannot measure them will be tempted to monitor conduct instead of output, which is precisely the overreach that discredited Ford's Sociological Department.

The practical conclusion is that a modern settlement has to be designed rather than inherited, and that its components are identifiable now: a revocable premium paid for demonstrated verification rather than assisted throughput; catch rates measured against a seeded baseline rather than asserted; unassisted work for junior staff budgeted as a training cost; and the resulting corrections captured as proprietary data. Settlements are already beginning to arrive in the AI economy through negotiation and regulation. Ford's real advantage in 1914 was not the money he spent. It was that he still had the initiative.

---

[1] Daniel M. G. Raff and Lawrence H. Summers, "Did Henry Ford Pay Efficiency Wages?," NBER Working Paper No. 2101 (1986), 9 and Tables 7 and 8.

[2] Nestor Maslej et al., Artificial Intelligence Index Report 2025 (Stanford Institute for Human-Centered Artificial Intelligence, 2025), chap. 2, "Technical Performance."

1. Introduction: why the analogy works

Comparisons between artificial intelligence and earlier industrial revolutions are common enough to have become suspect. The electric motor, the steam engine, the personal computer, and the assembly line are all invoked routinely, and the invocation is usually decorative: a historical anecdote placed at the front of an argument that would proceed identically without it. This report takes the comparison more seriously than that, and accepts the corresponding obligation to say what would make it fail.

The case for Fordism specifically rests on a structural similarity rather than a thematic one. Both technologies act on the production of something previously made by skilled individuals, and both act in the same way: by decomposing the work into standardised steps that a machine can carry between. In each case the immediate consequence is a collapse in unit cost, and the deeper consequence is that the scarce input changes identity. That second movement is what makes the Ford case worth studying now, because it is the part firms consistently miss while it is happening, and the part the historical record allows us to observe from beginning to end.

Ford's own record is unusually well documented for this purpose. The production data, the price series, the turnover and absenteeism figures, and the terms of the five-dollar day are all available, and the case has been examined closely enough by economic historians that its harder questions have been argued out.[3] That matters because the argument this report makes depends on mechanism rather than on outcome. It is not enough to observe that Ford paid well and prospered; the question is what the money bought and why the new production system made it worth buying.

Three questions organise what follows. What did the new technology make abundant, and what did it leave scarce? What instability did the production system generate that it could not itself absorb? And what did management have to buy, on what terms, in order to keep operating?

These are asked of Ford in Section 2, translated into their AI-era equivalents in Section 3, tested against the current empirical evidence in Section 4, and converted into practical guidance in Section 5.

The hypothesis that emerges can be stated in advance, since the report is more useful read against a clear claim than as a series of findings withheld until the end. It has three parts. The binding constraint in AI-mediated work has shifted from output volume to verification capacity, which means firms are currently paying for the thing their technology now supplies cheaply and underpaying for the thing it cannot supply at all. The advantage available from AI is not the cost decline, which is competed away industry-wide, but the proprietary residue of human correction, which means the same activity that secures quality is the activity that builds the moat. And no firm has yet made a settlement aimed at either, which places the AI economy in the position Ford occupied in late 1913: after the production revolution, before the response.

The analogy also has to be allowed to break, and the places where it does are reported rather than smoothed over. Ford owned his cost decline and could therefore share it; the AI cost decline is supplied by vendors and may be competed away before any firm captures it, which is a serious challenge to the whole argument and is confronted directly in Section 5.6. Ford's instability was measurable within weeks; its modern counterpart is not. Ford lowered the skill the task demanded, whereas AI appears to raise the capability of the worker, which is not the same mechanism even where the effect on returns to experience looks similar. In several of these cases the divergence carries more managerial value than the parallel does, which is the strongest argument that the comparison is doing real work rather than ornamental work.

[3] Shari Eli, Joshua K. Hausman, and Paul W. Rhode, "The Model T," Journal of Economic History 85, no. 1 (March 2025): 110–51. See also Raff and Summers, "Efficiency Wages."

2. The production revolution: how the mechanisms of Fordism worked

2.1 How the line changed the work and the market

The Model T revolutionised the automotive industry. Ford introduced the moving assembly line at its Highland Park plant in 1913, and the result was a collapse in cost and an expansion of the market at the same time.4

The fall in cost shows up most clearly in labour time. In the summer of 1913, before the line was installed, assembling a Model T chassis took 12.5 hours of labour; by December 1913, with the chassis moving continuously past fixed stations, it took 2 hours and 38 minutes.[5] Two factors drove the change: workers focused on a single narrow task, and the car itself was highly standardised, sold with a minimum of unnecessary features and, from 1914 to 1925, only in its famous black.[6] Lower cost fed directly into price. The Model T's real price halved between 1909 and 1915 and halved again between 1915 and 1921.[7] Plotted against cumulative production on logarithmic scales, the descent traces a near-straight line. This is the pattern later named the experience curve, and the Model T became its canonical illustration.[8] The result was a price-for-quality ratio that no competitor matched until the mid-1920s.[9]

The falling price expanded the market rather than merely capturing it. In 1909 the car was a luxury good, comparable in rarity and purpose to yachts today; by 1919 it was a mass-market good, and the growth came disproportionately from rural America. After mass production began, car registrations grew fastest in the most rural states, and by 1922 the Model T's share of cars in a county rose systematically with the rural share of its population.10 A machine that had signalled wealth was now within reach of the middle class.

The scale of Ford's advantage was stark. In 1914, Ford produced 260,722 cars with 13,000 employees, while all other US manufacturers combined produced 286,770 with 66,350.[11] But the same line that collapsed the cost of a car also hollowed out the skill of building one, and the consequences arrived quickly.

2.2 What the line did to labour

The assembly line changed both the nature of the work and the workforce doing it. Craft skill gave way to narrow, repetitive tasks set by the machine, and the instability that followed became a problem management could not ignore.

Worker dissatisfaction at Ford had three distinct sources: the task, the conditions, and the treatment. The task was monotonous and required almost no skill; a Yale student who spent a summer at Highland Park observed that most jobs could be mastered in a matter of minutes.[12] Rest was scarce, because the line set the pace, unlike in traditional

[4] Eli, Hausman, and Rhode, "The Model T," 110–11 and 128–30. Their price and production series draw on Bruce W. McCalley, Model T Ford: The Car That Changed the World (Iola, WI: Krause Publications, 1994).

[5] Allan Nevins, Ford: The Times, the Man, the Company (New York: Charles Scribner's Sons, 1954), 473, cited in Eli, Hausman, and Rhode, "The Model T," 129.

[6] Eli, Hausman, and Rhode, "The Model T," 128.

[7] Eli, Hausman, and Rhode, "The Model T," 129–30.

[8] Bruce D. Henderson, "The Experience Curve," BCG Perspectives (Boston: The Boston Consulting Group, 1968); The Boston Consulting Group, Perspectives on Experience (Boston, 1972); William J. Abernathy and Kenneth Wayne, "Limits of the Learning Curve," Harvard Business Review 52, no. 5 (1974): 109–19.

[9] Daniel M. G. Raff, "Making Cars and Making Money in the Interwar Automobile Industry," Business History Review 65, no. 4 (1991): 721–53.

[10] Eli, Hausman, and Rhode, "The Model T," 122–38.

[11] Nevins, Ford, 488, cited in Eli, Hausman, and Rhode, "The Model T," 129.

[12] Raff and Summers, "Efficiency Wages," 8 (Yale engineering student on the mastery of movements).

manufacturing where the worker controlled it. The conditions were poor as well: workdays were long, pay was low, and housing was inadequate. And the treatment was arbitrary, with foremen holding largely unaccountable power over hiring, firing, and pay.[13]

Altogether this produced staggering churn. In 1913 annual turnover reached 370 percent, and Ford had to hire over 50,000 men to maintain an average workforce of fewer than 14,000.[14] Much of this was not formal resignation: of the roughly 7,300 workers who left in March 1913, about 71 percent were "five-day men" who had simply stopped showing up and, after five straight unexcused absences, were deemed to have quit. On top of this came a daily absenteeism rate of 10 percent, which forced Ford to run the plant with more than a thousand inexperienced stand-ins on an average day.[15]

In a craft shop, this churn would have been costly but survivable. On the assembly line, where each station depended on the one before it, a single missing worker could stall the whole line. The production system that had caused the instability was also the system least able to tolerate it. Something had to change.

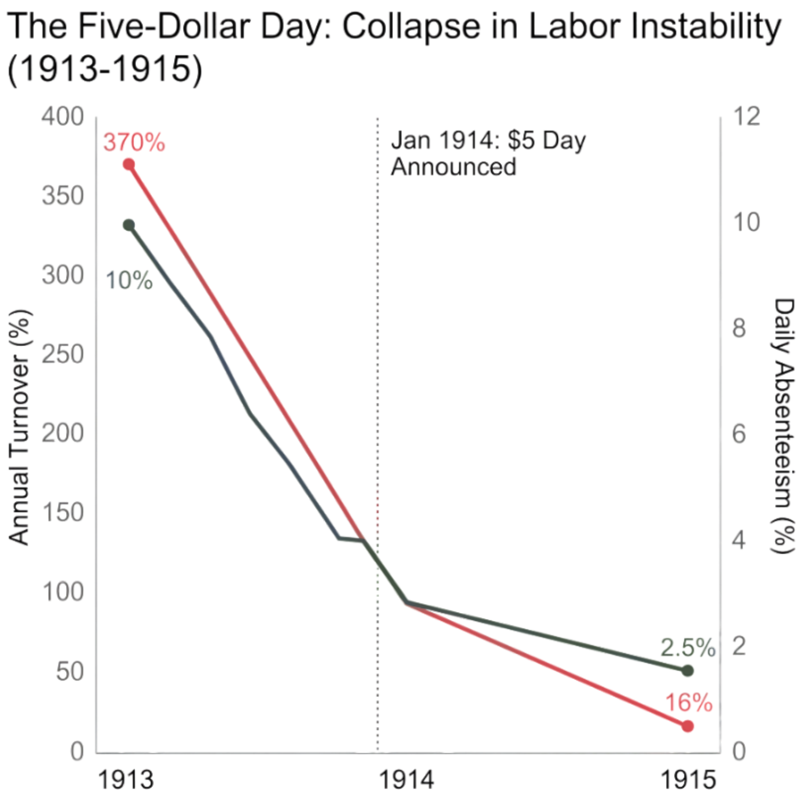

Figure 1. The five-dollar day and the collapse in labour instability. Turnover and absenteeism fell sharply after Ford changed the terms of employment.

2.3 The five-dollar day

Ford's answer came on 5 January 1914: the five-dollar day, which more than doubled pay for eligible workers and cut the working day from nine hours to eight. The announcement made Ford's name known to virtually every newspaper reader in the United States.[16] The headline figure, however, concealed the structure of the offer. The five dollars was not a straightforward wage increase but a base wage combined with a conditional profit-sharing bonus, each carrying very different terms.

The base was the guaranteed portion, and it changed comparatively little: a worker earning

[13] Raff and Summers, "Efficiency Wages," 12–13, on John R. Lee's 1913 report on the causes of dissatisfaction.

[14] Raff and Summers, "Efficiency Wages," 9 and Table 7, drawing on Sumner Slichter, The Turnover of Factory Labor (1921), 244.

[15] Raff and Summers, "Efficiency Wages," 9 and Table 8.

[16] Daniel M. G. Raff, "Buying the Peace: Wage Determination Theory, Mass Production, and the Five-Dollar Day at Ford" (revised MIT PhD thesis, 1997), 2, quoted in Eli, Hausman, and Rhode, "The Model T," 133.

$2.34 a day under the old schedule continued to earn roughly that under the new plan.[17] What changed alongside it was the eight-hour day, run across three shifts, which allowed Ford to keep the plant working around the clock. The base asked nothing of the worker in return.

The bonus was where the money and the conditions sat. Ford paid a further $2.66 a day, but only to workers who met the company's standards for how they lived outside the factory. The newly created Sociological Department sent investigators into workers' homes to assess eligibility: qualifying required abstaining from alcohol, refraining from abusing one's family, taking in no boarders, keeping a clean home, and contributing regularly to savings.[18] Only male workers qualified at the outset. For a largely immigrant, blue-collar workforce, the scheme raised living standards materially, but it did so through an intrusive paternalism that made pay conditional on conformity in workers' private lives.

The puzzle is why Ford did it at all. Thousands queued outside Highland Park when the announcement broke, so labour was plainly available at a fraction of the cost. No market pressure forced a profit-maximising firm to pay double the going rate. What Ford got for the money, and what that reveals about his new production system, is the subject of the next section.

2.4 A system of deliberate choices

The results were immediate. Turnover collapsed from 370 percent in 1913 to 54 percent in 1914 and 16 percent by 1915, and absenteeism fell from 10 percent of the workforce on a given day to 2.5 percent within a year of the announcement.[19] Ford later described the five-dollar day as one of the finest cost-cutting moves the company ever made.[20] The question is what, mechanically, the money bought.

The intuitive answer, that higher pay simply induced men to work harder, does not survive contact with the assembly line. The line, not the worker, set the pace, and a worker could not meaningfully choose to produce more. What he could choose was whether to show up, whether to stay, and whether to accept a discipline he had previously walked away from. This is not to say the work did not intensify. The production foreman William Klann later recalled management reasoning that since the men were receiving twice the wages, the company wanted twice the work, and simply raising the speed of the lines.[21] The distinction matters. Ford did not buy harder work by inducing it; he

bought the ability to impose it. A workforce turning over at 370 percent a year could not have been driven faster without emptying the plant, and the bonus, conditional and revocable, made staying worth the intensification.

This matters because the production revolution had changed what management was buying. On the craft floor, output depended on how much an individual could produce. On the line, output depended on whether the line kept moving, which meant it depended on whether the same men returned to the same stations every morning. Retention, not raw effort, became the binding constraint, and Ford's wage was aimed precisely at it. This reading sits close to the conclusion Raff and Summers themselves reach: the five-dollar day fits the broad logic of efficiency wages, but not the textbook version in which higher pay deters shirking. They attribute the gain instead to the rising cost of shirking under interdependent production and to rent-sharing, or in their phrase, to buying the peace.[22] On Ford's line there was little scope to shirk. There was enormous scope to leave.

The evidence suggests the wage worked. Ford's net income rose from $13.5 million in 1912 to $27.1 million in 1913, $31.8 million in 1914, and $40.3 million in 1915, through the very period in which he doubled his wage bill.[23] Two qualifications belong here. Raff and Summers calculate that reduced turnover and training costs offset only a

[17] Raff and Summers, "Efficiency Wages," 16.

[18] Raff and Summers, "Efficiency Wages," 16–18. The eligibility conditions were age (men over 22), six months' service, and certification by the Sociological Department.

[19] Raff and Summers, "Efficiency Wages," Tables 7 and 8. The authors caution at 28 that the Detroit downturn of late 1913 may account for a substantial share of the fall in turnover.

[20] Henry Ford and Samuel Crowther, My Life and Work (New York, 1922), 126, 127, 147, quoted in Raff and Summers, "Efficiency Wages," 3–4.

[21] Raff and Summers, "Efficiency Wages," 30, quoting the production foreman William Klann's recollection that management raised line speed after the wage increase.

[22] Raff and Summers, "Efficiency Wages," 4 and 29–32.

[23] Raff and Summers, "Efficiency Wages," Table 2 (Ford Motor Company net income, 1910–1915). Nominal figures; the table also reports real values in 1910 dollars.

small fraction of the programme's cost, so the wage did not pay for itself in those terms alone.[24] And rival motives exist, including the enormous publicity the announcement generated and the value of forestalling union organisation.[25] None of these is incompatible with the economics.

Whichever weight one assigns, the underlying logic is the same: Ford shared the surplus his new technology had created because he could not run that technology at full intensity with a workforce that would not stay. This was strategy, not charity.

The lesson sits in the conditionality rather than the generosity. Ford attached the money to the behaviours his new system actually required, and he could take it away. Where he overreached was in the reach of that condition, extending it from the factory floor into workers' homes. The conditionality, not the intrusion, is what transfers, and it is the principle that Section 5 will return to.

[24] Raff and Summers, "Efficiency Wages," 29–31. Reduced turnover and training costs offset only a small fraction of the programme's cost on the authors' calculations.

[25] Raff and Summers, "Efficiency Wages," 15–16 on publicity, and 32 on "buying the peace" and the avoidance of collective output restriction.

3. Levers of equivalence

Section 2 traced a single argument through the Fordist case: a new production technology collapsed costs, expanded a market, remade the work, generated an instability the system could not absorb, and forced management to identify what had become scarce and pay for it. That sequence is the spine of this section as well. Each of the seven levers below asks the same question of AI-mediated work that Ford eventually had to answer at Highland Park, and the levers build toward it: what does this technology make abundant, what does it leave scarce, and what must a firm therefore buy? A comparison of this kind earns its place only if the mappings are structural rather than decorative, so each lever is tested on whether the same underlying mechanism is at work and whether it survives the points at which the two eras come apart. Those divergences are reported rather than smoothed over, and several of them carry more managerial value than the parallels do.

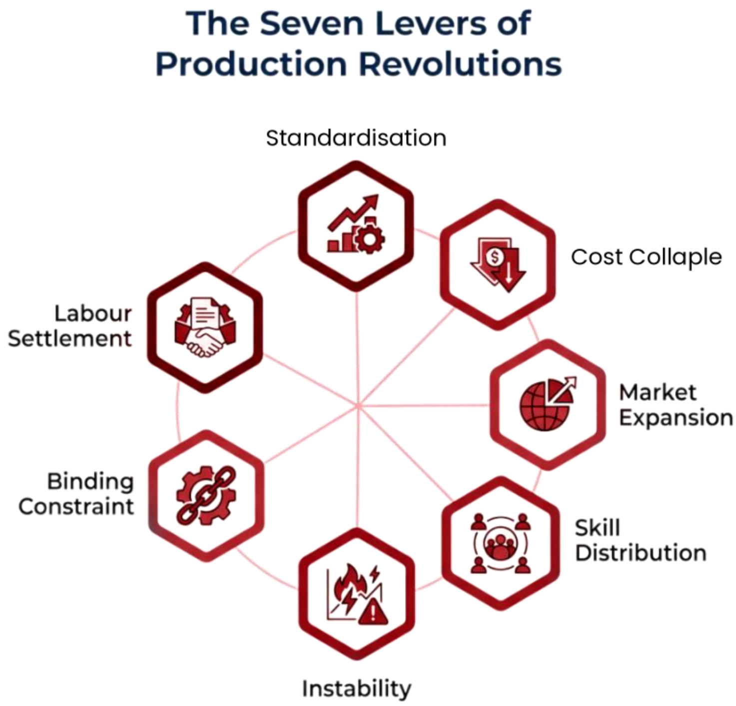

Figure 2. The seven levers of the production-revolution analogy. The framework translates the Fordist mechanism into an AI-era sequence from standardisation to a prospective settlement.

The seven levers are not all the same kind of thing, and the difference matters for what follows. Five of them describe conditions the technology imposes on a firm whether or not it welcomes them: standardisation, cost collapse, market expansion, the reshaping of skill, and the instability that results. The sixth is a diagnosis rather than a condition, naming what has become scarce. Only the seventh is a decision, and it is the one no firm has yet made. Section 5 therefore has far less discretion available to it than a list of seven might suggest. The seven are discussed in turn below.

3.1 Standardisation

Standardisation is visible in both cases, and divides into three parts: product, process, and pace. Ford's one model in one colour finds its counterpart in the convergence of model-assisted output, which is measurable: after AI was deployed in one large customer-support operation, textual similarity between high- and low-skill agents rose from 0.55 to 0.61, with the lower-skilled changing most, so that they came to write more like the strongest performers.[26] Ford chose his standardisation, and could have traded it against variety at a cost; convergence under AI arrives whether or not a firm wants it, which is the first hint that this technology sets terms rather than merely offering options. Work at Highland Park consisted of single, narrow tasks; AI-mediated work consists of

[26] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work," Quarterly Journal of Economics 140, no. 2 (2025): 889–942, at 929–30, measuring cosine similarity between the text of high- and low-skill agents before and after deployment.

small, model-mediated steps, and the measured gains follow the same pattern of narrow tasks completed faster.[27] Pace is where the two eras separate. The line dictated the speed of work at the plant and the worker could not choose otherwise, whereas today's assistants are consulted at the worker's discretion. This is a real disanalogy, but it may be a temporary one. Most firms currently use AI the way early factories used electric motors, as a faster tool inside an unchanged workflow, rather than redesigning the workflow around it. The shift to agentic systems, which execute multi-step workflows independently, is the point at which pace stops being voluntary.[28] On this reading, most organisations in 2026 have not yet had their Highland Park moment, which means the levers that follow describe a transition still in progress rather than one already complete.

3.2 Cost collapse

Costs fell steeply in both eras. At Ford, labour time per chassis collapsed as the process was redesigned around the line, and unit cost fell as cumulative volume grew. In the AI era, the cost of a fixed level of performance has fallen by more than two orders of magnitude in under two years.[29] The direction is the same; the mechanism is not, and the difference is where the managerial value sits. Ford's costs fell because of what Ford did. AI costs fall because of what the model vendors do.

Ford's cost decline was proprietary. His cumulative volume drove his costs down, no competitor shared in the benefit, and the resulting price-for-quality advantage held for roughly a decade.[30]

The AI cost decline is supplied. It arrives from a small number of vendors and reaches every competitor on the same day. Even the choice between paid and freely available models has narrowed as a source of raw capability: the gap between leading closed-weight and open-weight models fell from 8.0 percent to 1.7 percent on benchmark tasks between January 2024 and February 2025, though benchmark parity is not the same as equivalent business value.[31] Ford's curve was a moat; the AI cost curve is closer to a commons. The implication is uncomfortable for any firm treating cheap inference as a source of advantage: falling costs pass through to prices across the industry and leave relative position unchanged. A firm's own experience curve has to be built out of what cannot be bought from the same supplier: proprietary workflows, proprietary data, and feedback loops in which its own people's corrections improve its own system. Each of those depends on human judgement applied to model output, which is the constraint the later levers converge on.

3.3 Market expansion

Falling price expanded the market in Ford's case, and the same mechanism is visible now, though the two are easily conflated with something else. The real price of the Model T halved twice in twelve years, and the growth that followed came disproportionately from rural America rather than from the wealthy urban buyers who had bought cars before.[32] The market for the product expanded because the product got cheap, not because more manufacturers installed assembly lines. The AI-era claim has to be framed the same way, in terms of what a buyer of cognitive output now pays rather than how many firms have adopted the technology. On that measure the movement is dramatic: querying a model performing at GPT-3.5 level fell from $20.00 to $0.07 per million tokens between November 2022 and October 2024.[33] At that price, output that once required a scarce and expensive expert is within reach of buyers who could not previously commission it at all, which is the same luxury-to-mass movement Ford's price cuts produced.

One clarification keeps this lever honest. Organisational adoption figures, often cited here, measure diffusion of the production technology across firms, which is the analogue of counting how many manufacturers installed an

---

27 Brynjolfsson, Li, and Raymond, "Generative AI at Work," 889 and 907.

28 Marie-Christine Fregin, Danique Eijkenboom, Pelin Özgül, Mantej Pardesi, and Nicholas Rounding, "Artificial Intelligence in Firms: A Synthesis of Micro-Level Evidence on Productivity, Work, and Skills," ROA Policy Brief No. 2026/2E (Maastricht University, 2026), 6–7, on compound AI agents that independently execute entire task packages.

29 Maslej et al., AI Index Report 2025, chap. 2.

30 Raff, "Making Cars and Making Money," 721–53.

31 Maslej et al., AI Index Report 2025, chap. 2. The gap is measured on benchmark tasks and does not by itself indicate equivalent business value.

32 Eli, Hausman, and Rhode, "The Model T," 122–38.

33 Maslej et al., AI Index Report 2025, chap. 2. The cost of querying a model scoring at GPT-3.5 level (64.8 percent on MMLU) fell from $20.00 per million tokens in November 2022 to $0.07 by October 2024.

assembly line rather than asking who could now afford a car; the incidence question for AI is not yet settled the way registration data settles it for Ford. What is settled is the direction, and it points at the same consequence in both eras: when the cost of producing something collapses, the constraint on the business stops being the cost of producing it.

3.4 Skill and the shape of expertise

Both technologies reshaped the distribution of skill, but not by the same mechanism, and the difference is worth stating precisely. It helps to separate three things: the ceiling, meaning what the task demands of whoever performs it; the floor, meaning what a given worker can produce; and the slope, meaning how much additional experience is worth.

Ford lowered the ceiling. The line stripped skill out of the task itself, so that jobs which had required a craftsman could be mastered in minutes, and the worker's own capability was left untouched while the task fell to meet it. AI does something close to the opposite: it raises the floor. Models capture the problem-solving patterns of strong performers and make them available to everyone, so the lowest-skilled workers gain around 36 percent in resolution speed while the highest-skilled gain nothing measurable.[34] What the two eras share is the effect on the slope. Both compress the returns to experience, and in the AI case the compression is dramatic: two months of tenure with model assistance matches roughly six months without it.[35] As codified, transferable knowledge becomes cheap to reproduce, the premium shifts to tacit judgement and to the kind of complex decision-making a model cannot yet supply.[36] The levelling is real but not yet universal: synthesising the firm-level evidence, Fregin and colleagues find substitution still dominant and note that where tasks turn on tacit knowledge the effect can invert, with novices following suggestions mechanically while experts use the same tool to sharpen judgement they already hold.[37] That inversion is itself informative, because it locates precisely where human contribution still decides the outcome. It also raises a question Ford never had to face. If novices reach veteran output without accumulating the experience that produced veterans, the pipeline generating future expertise is left unexplained, and the same synthesis names this erosion as a leading risk: when codifiable work is taken over wholesale, the intermediate stages through which judgement is built disappear.[38] A firm that depends on judgement must therefore produce it deliberately, which is a cost it did not previously carry.

3.5 Instability

Every production system generates its own characteristic failure, and both eras illustrate the point, though the failures are opposites. Ford's instability was rejection. Workers left, at 370 percent a year, or simply did not appear, at 10 percent on a given day, and the line could not absorb either; both fell sharply once the wage changed, turnover to 54 percent within a year and 16 percent within two, and absenteeism to 2.5 percent.[39] The AI era's instability is the reverse: workers accept the system too readily. Over-reliance on model output degrades performance in complex decision-making and contributes to long-term cognitive deskilling,[40] and plausible but hollow output proliferates wherever nobody is deliberately checking it.[41] Ford's problem was that his workers walked away from the system; the modern problem is that they lean on it.

A second difference matters for management. Ford could see his instability every morning in the form of empty stations and unfilled shifts, which is why he could measure it, price it, and act on it. Cognitive deskilling and unchecked output produce no equivalent signal. They appear as work delivered on time, and the deterioration

---

[34] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 911.

[35] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 916 and Figure V.

[36] Fregin et al., "Artificial Intelligence in Firms," 5–6, distinguishing substitution, levelling, and expertise enhancement.

[37] Fregin et al., "Artificial Intelligence in Firms," 6, reporting substitution as the dominant mechanism to date and noting that broad-based levelling remains absent.

[38] Fregin et al., "Artificial Intelligence in Firms," 8, listing erosion of the skills pipeline among the leading risks.

[39] Raff and Summers, "Efficiency Wages," 16, 29–31, and Tables 7 and 8.

[40] Fregin et al., "Artificial Intelligence in Firms," 5.

[41] Laura Schneider and Maria Kharlamova, "Mind over Machine: Navigating Human Metacognition When Using Generative AI," short paper, Proceedings of the Forty-Sixth International Conference on Information Systems (Nashville, 2025), 1–3; Fregin et al., "Artificial Intelligence in Firms," 4, drawing on Kate Niederhoffer et al., "AI-Generated 'Workslop' Is Destroying Productivity," Harvard Business Review, 22 September 2025.

surfaces only later and indirectly. Ford could price his instability and act on it; this one has to be looked for before it can be managed at all. Mitigating it depends on deliberate work design and on prompting critical evaluation rather than passive acceptance,[42] which is to say it depends on things a firm must choose to build, because nothing in the ordinary run of operations will reveal the problem on its own.

3.6 The binding constraint

This is the argument the two sections turn on. Section 2 established that the assembly line changed what management was buying: on the craft floor, output depended on individual effort, but on the line the pace was fixed by management rather than by the worker, and output depended on whether the same men returned to the same stations every morning, which made retention rather than autonomous effort the binding constraint.[43] The same question, asked of AI-mediated work, produces a different answer again. Output volume is no longer scarce, because the model supplies as much of it as is wanted at a collapsing price. What remains scarce is the judgement that decides whether a given output is any good, and the human role shifts accordingly from production toward evaluation and exception handling.[44] Effort, then retention, and now verification capacity: in each regime the constraint moved to whatever the technology could not supply for itself, and in each case management had to work out what it was actually buying before it could pay for it correctly.

This closes the circle opened in 3.2. The advantage that cheap inference cannot confer has to be built from proprietary workflows, proprietary data, and feedback loops, and each of those is produced by the same activity that the binding constraint identifies: a person checking model output and correcting it. The correction is not merely a quality control cost. It is the raw material of the only advantage a competitor cannot buy from the same vendor, because the competitor is buying the model and not the record of how one firm's people improved it. Stated as a proposition for management, verification is not an overhead on AI adoption; it is the mechanism by which AI adoption becomes an advantage at all, and a firm that under-resources it is not simply risking quality but declining to build the asset.

3.7 The settlement that has not happened

The seventh lever is the one with an empty cell, and the emptiness is the point. Ford answered his binding constraint by sharing the surplus his technology had created, conditionally and revocably, in a form aimed precisely at the behaviour his system required.[45] No comparable settlement has been made in the AI era. That claim needs defending, because plausible candidates exist: the 2023 Writers Guild agreement placed limits on studio use of AI, several large employers have announced retraining commitments or no-layoff pledges, and AI usage policies are now common. None of these is a settlement in the sense Ford's was.

Four features made the five-dollar day what it was. It transferred a substantial share of the surplus the technology had created. It was conditional, so the money was tied to something the firm needed in return. It was revocable, so the condition had force. And it was aimed precisely at the binding constraint the new system had produced, rather than at instability in general. Measured against those four, the current examples fall short in different ways. Usage policies impose conditions but transfer nothing. Retraining pledges transfer something but are rarely conditional or revocable, and are seldom aimed at verification capacity specifically. The Writers Guild agreement is the closest, but it restricts the technology rather than sharing its surplus, and it addresses displacement rather than the constraint identified in 3.6.

A fifth feature separates Ford's case more sharply still, and it is the one that matters most for what follows. Ford acted unilaterally and pre-emptively. Nobody struck, legislated, or bargained him into the five-dollar day; he identified what his system needed and paid for it before the choice was taken out of his hands. Most of the current candidates are the opposite: negotiated under the threat of a strike, or written in response to regulation. Settlements, in other words, have already begun to arrive in the AI economy. They are simply being imposed rather than designed. The AI economy in 2026 stands roughly where Ford stood in late 1913, with one difference that should concern any firm reading this: Ford still had the initiative.

---

[42] Schneider and Kharlamova, "Mind over Machine," 5 and Table 1: in a pilot experiment (n = 60), raising the visual salience of a standard error disclaimer significantly increased response times and self-regulatory actions, though it did not significantly improve the calibration between confidence and actual performance.

[43] Raff and Summers, "Efficiency Wages," 16, 29–31, and Tables 7 and 8.

[44] Fregin et al., "Artificial Intelligence in Firms," 5, on monitoring and validation roles as cognitively asymmetric; Schneider and Kharlamova, "Mind over Machine," 1–3, on the metacognitive demand this places on the user.

[45] Raff and Summers, "Efficiency Wages," 16–18 and 32.

Read this way, the Fordist case supplies not only a precedent but a warning about how such a settlement should be designed. Ford tied his money to the behaviours the new system needed, which was the sound part, and then extended the condition from the factory floor into workers' homes, which was not. The modern equivalent of that overreach is not hard to identify: monitoring how employees think and work with these systems rather than what they produce with them. The distinction between conditioning on output and conditioning on conduct is the one Ford failed to hold, and it is the substance of the lessons and counter-lessons that follow in Section 5.

4. The AI evidence

Section 3 set out seven levers. This section presents the empirical evidence behind each in the same order, so that the translation and the data supporting it can be read against one another. The evidence base is thinner than the Fordist record in one specific respect worth stating at the outset: the Ford case is closed and can be assessed with a century of hindsight, whereas the AI evidence consists largely of field experiments conducted within three years of deployment, in a small number of occupations, with wide confidence intervals. Where the findings are contested or the intervals are wide, this is noted.

4.1 Standardisation and productivity

The clearest productivity evidence comes from two field settings. In a large customer-support operation, access to a generative AI assistant raised issues resolved per hour by roughly 15 percent.[46] In software development, three randomised field experiments across Microsoft, Accenture, and an anonymous Fortune 100 manufacturer found a 26 percent increase in completed tasks among the developers who took up the tool, though the standard error of 10.3 on that estimate means the true effect could plausibly be half or well over that figure.[47] The mechanism in the support setting decomposes into shorter handling times, more issues handled concurrently, and a modest rise in resolution rates,[48] which is recognisably the same pattern of narrow tasks completed faster that the line produced at Highland Park.

4.2 Cost collapse

The cost of a fixed level of model performance has fallen faster than almost any comparable input in industrial history. Querying a model scoring at GPT-3.5 level fell from $20.00 per million tokens in November 2022 to $0.07 by October 2024, and depending on the task, inference prices have fallen between 9 and 900 times per year.[49] The competitive significance of this, as Section 3.2 argued, lies less in the magnitude than in who captures it. The narrowing of the leading closed-weight to open-weight performance gap from 8.0 percent to 1.7 percent within roughly a year indicates how little raw model access now separates firms.[50]

4.3 Market expansion

The price evidence above is also the market-expansion evidence, since a collapse in what a buyer pays for cognitive output is precisely the mechanism that brings new buyers into the market. What the current data does not establish is incidence. Organisational adoption has risen sharply, from 55 percent to 78 percent of organisations in a single year,[51] but as Section 3.3 noted, this measures diffusion of the production technology across firms rather than which buyers are newly served by its output. The distributional question that registration data answers for the Model T has no equivalent answer here yet.

---

[46] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 907. The preferred specification gives a 15.2 percent increase in issues resolved per hour.

[47] Zheyuan (Kevin) Cui, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz, "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers," Management Science (2025), doi 10.1287/mnsc.2025.00535, 1 and Table 3. The pooled estimate is a 26.08 percent increase in weekly completed pull requests (SE 10.3) across 4,867 developers. It is a local average treatment effect for those who adopted the tool, estimated by instrumenting adoption with random assignment; the trials ran in 2022–23 on the version of GitHub Copilot then available.

[48] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 907–11, decomposing the gain into reduced handling time, higher concurrency, and a modest rise in resolution rates.

[48] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 907–11, decomposing the gain into reduced handling time, higher concurrency, and a modest rise in resolution rates.

[49] Maslej et al., AI Index Report 2025, chap. 2.

[50] Maslej et al., AI Index Report 2025, chap. 2.

[51] Maslej et al., AI Index Report 2025, chap. 4, "Economy."

4.4 Skill and expertise

The distributional finding is the most striking in the literature and the most easily overstated. In the customer-support setting, the lowest-skilled workers gained around 36 percent in issues resolved per hour while the highest-skilled gained nothing measurable and showed a slight decline in quality.[52] Workers with two months of tenure and model assistance matched the performance of workers with over six months and none, which the authors describe as moving more quickly down the experience curve.[53] The software experiments point the same way: in the one trial where developer characteristics were observable, short-tenure and junior developers both took up the tool more readily and gained more from it, though the gap between them and their senior colleagues is not statistically significant.[54]

Two qualifications matter. The system in the support study was fine-tuned on the conversations of high-performing agents, so what it diffused was that firm's own best practice rather than a general property of the technology. And the broader firm-level synthesis finds substitution still dominant, with broad-based levelling not yet materialised and the effect capable of inverting

where tasks turn on tacit knowledge.[55] The levelling result is real and well identified; it is not yet a general law.

Figure 3. AI's levelling effect on productivity. The productivity gains are concentrated among lower-skilled workers in the cited customer-support evidence.

4.5 Instability

The evidence on the AI era's characteristic failure is weaker than the Fordist turnover record, and honest reporting requires saying so. What exists is consistent. Over-reliance on model output degrades performance in complex decision-making and is associated with long-term cognitive deskilling, and the erosion of the skills pipeline is

[52] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 911. The system was fine-tuned on the conversations of high-performing agents, so the levelling observed is a property of that deployment rather than of AI in general.

[53] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 916 and Figure V.

[54] Cui et al., "Effects of Generative AI on High-Skilled Work," sec. 5 and fig. 6. Developer characteristics were observable only in the Microsoft trial: short-tenure developers were 9.4 percentage points more likely to adopt (81.6 against 72.1 percent) and raised output by 27 to 39 percent against 8 to 13 percent for long-tenure developers, while junior developers gained 21 to 40 percent against 7 to 16 percent for senior ones. The authors stress that these differences are not statistically significant at conventional levels.

[55] Fregin et al., "Artificial Intelligence in Firms," 6.

identified as a leading organisational risk.[56] Users overestimate their own performance when model assistance is available and accept suggestions without adequate critical evaluation.[57] Against this, the exit-type instability that dominated Ford's crisis appears to move in the opposite direction: attrition fell among agents with under six months of tenure who had model access.[58] This supports the inversion argued in Section 3.5, in which the modern failure is over-acceptance rather than rejection, but it should be read as a documented risk rather than as a measured operational crisis of the kind Ford faced.

4.6 The binding constraint

No study measures verification capacity directly, and the constraint is inferred from converging evidence rather than demonstrated. Three findings support the inference. Gains accrue where humans review rather than produce, since the productive step is increasingly supplied by the model. Monitoring and validation work is cognitively asymmetric, consisting of infrequent high-stakes judgements with limited feedback, which is a poor environment for maintaining skill.[59] And interventions that raise deliberation do not automatically improve accuracy: raising the visual salience of an error disclaimer significantly increased response times and self-regulatory actions in a pilot experiment, but did not significantly improve the calibration between confidence and actual performance.[60] That last result is the most consequential for Section 5, because it indicates that verification quality cannot be secured by prompting alone.

4.7 The absent settlement

The seventh lever has no evidence because there is nothing yet to measure, which is the finding rather than a gap in the literature. The nearest thing to a proposal in the empirical record comes from the customer-support study itself: its authors observe that top performers increased adherence to model suggestions even where this slightly degraded their own quality, which reduces the flow of novel human solutions available to train future systems, and they raise the possibility of compensating workers for contributions to model training.[61] That is a proto-settlement argument arising directly from the data, and Section 5 takes it up.

---

[56] Fregin et al., "Artificial Intelligence in Firms," 4–5 and 8.

[57] Schneider and Kharlamova, "Mind over Machine," 1–3; Fregin et al., "Artificial Intelligence in Firms," 4, drawing on Niederhoffer et al., "AI-Generated 'Workslop.'"

[58] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 920–21, reporting reduced attrition among agents with under six months of tenure.

[59] Fregin et al., "Artificial Intelligence in Firms," 5; Schneider and Kharlamova, "Mind over Machine," 1–3.

[60] Schneider and Kharlamova, "Mind over Machine," 5 and Table 1.

[61] Brynjolfsson, Li, and Raymond, "Generative AI at Work," 936–37, noting that adherence among top performers may reduce the flow of novel solutions available to train future models, and raising compensation for training contributions as a possible response.

5. The modern settlement: lessons and counter-lessons

Sections 3 and 4 established the levers and the evidence behind them, and together they produce a diagnosis: output volume is abundant, verification capacity is scarce, and no firm has yet made a settlement aimed at the constraint this creates. This section takes that diagnosis as given and asks the question a chief executive actually faces. If the settlement has to be designed rather than inherited, what does designing it involve, and what should a firm start doing about it?

Five lessons follow, then two counter-lessons. The first four lessons concern the settlement itself: what to pay for, what to protect, what to build into the work, and what to own. The fifth concerns timing. The counter-lessons are the two ways the Fordist precedent misleads, and the first of them is a challenge to the analogy on which this report rests.

5.1 Re-price the verification bottleneck

The binding constraint has moved from production to judgement, and compensation has not moved with it. Most firms still pay for throughput in roles where throughput is no longer the scarce input, and the effect is to reward exactly the behaviour the system produces too much of. Re-pricing this is the closest modern equivalent to the five-dollar day, and Ford's design gives the structure: share a real part of the surplus, attach a condition, keep the ability to withdraw it, and aim it at the actual constraint rather than at instability in general.[62] Working through those four in the AI case turns an observation into a design.

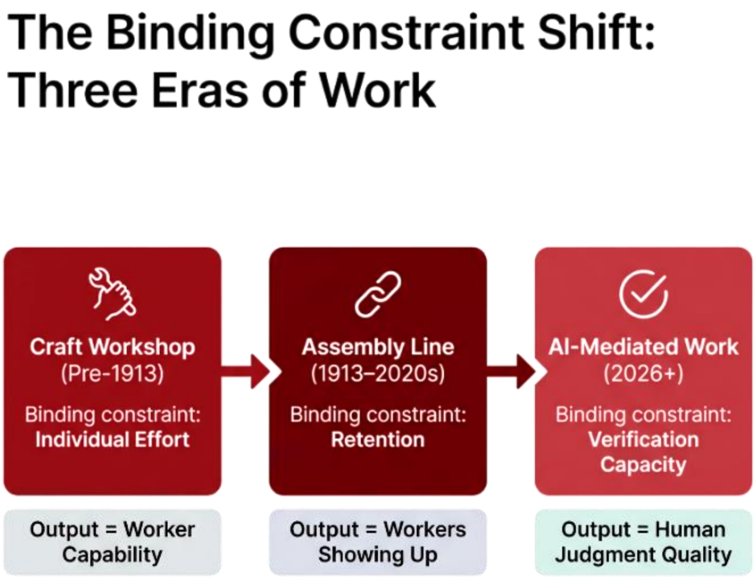

Source: Author's analysis drawing on Raff & Summers (1986) and Brynjolfsson et al. (2025)

Figure 4. The binding constraint shifts across production regimes. The scarce input moves from individual effort to retention and, in AI-mediated work, verification capacity.

What to reward. Not the volume of AI-assisted output, and not self-reported diligence, but the catch: the demonstrated identification of a model error, an unsupported claim, or an inapplicable recommendation before it reached a client or a decision. This is the measurable analogue of showing up at the same station every morning.

How to measure it. Verification is only distinguishable from approval if the work leaves a trace, which means the reviewing step has to record what changed. Two mechanisms are available without new technology: requiring a substantive edit log or rejection rationale on AI-assisted work before sign-off, and seeding a known error rate into

[62] Raff and Summers, "Efficiency Wages," 16–18 and 32.

review queues so that catch rates can be measured against a baseline rather than asserted. The second is the more important, because it converts an invisible failure into a monitored one and gives the firm the metric Ford already had in his turnover figures. The experimental caution from 4.6 applies directly here: raising the salience of an error warning increases deliberation without measurably improving calibration,[63] so an incentive must be tied to catches found rather than to care claimed.

When to withdraw it. Ford's bonus had force because it was revocable. A verification premium that becomes a permanent salary component stops conditioning anything within a year. Tying it to sustained measured catch rates, reviewed at a fixed interval, keeps the condition live.

One caveat belongs here rather than in a separate lesson. Compound agents that execute entire task packages are already being deployed, and if such systems eventually close the verification loop themselves, a settlement anchored on human verification would be anchored on a role in decline.[64] This is an argument for making the premium revocable and reviewable, which the design above already requires, rather than an argument for waiting.

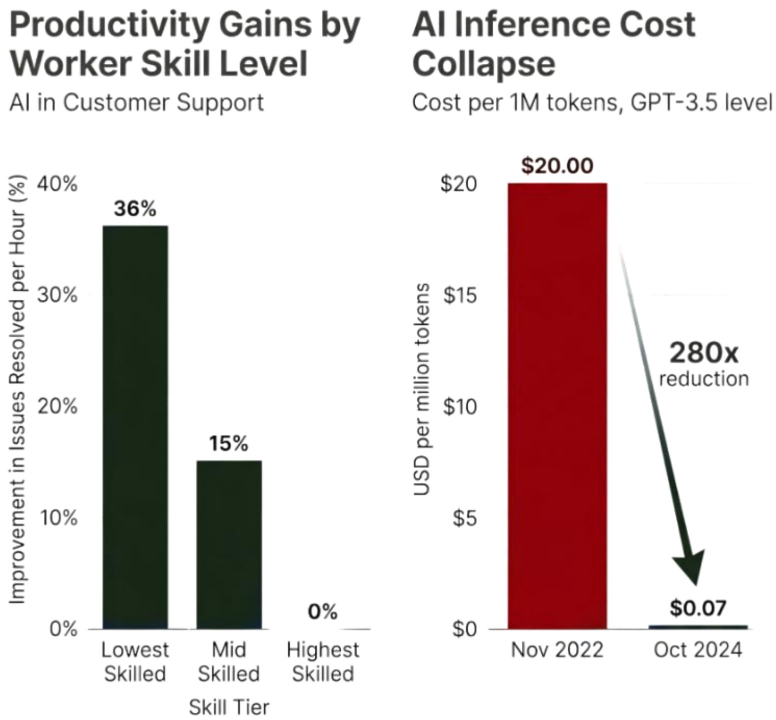

Figure 5. Productivity gains by worker skill and the collapse in AI inference cost. The evidence combines the distribution of productivity gains with the rapid decline in the cost of model inference.

5.2 Manufacture the expertise pipeline

Where AI levels performance, it does so by lifting the bottom rather than the top, and the compression of returns to experience is steep. The pipeline risk holds regardless of which mechanism dominates: if junior staff reach adequate output without passing through the intermediate stages that build judgement, the firm consumes expertise it is no longer producing.[65] Since 5.1 pays for judgement, and judgement is what stops being manufactured, this is the lesson that makes the first one sustainable.

Two practical moves follow. The first is to protect a deliberate share of unassisted work for junior staff, not as a productivity sacrifice but as training cost, budgeted and defended as such. The second is to design review allocation for exposure rather than efficiency: routing junior reviewers across a wide range of case types builds the pattern library that judgement rests on, whereas routing them to whatever is fastest to clear does not. Both are decisions about work allocation that a firm can make next quarter without new systems.

[63] Schneider and Kharlamova, "Mind over Machine," 5 and Table 1.

[64] Fregin et al., "Artificial Intelligence in Firms," 6–7, on compound AI agents that independently execute entire task packages.

[65] Fregin et al., "Artificial Intelligence in Firms," 8, listing erosion of the skills pipeline among the leading risks.

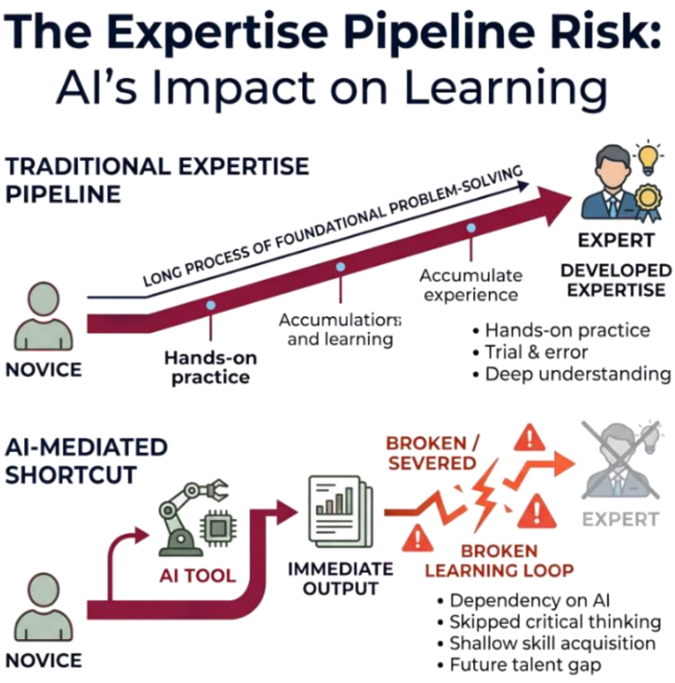

Figure 6. The expertise-pipeline risk created by AI-mediated work. When AI compresses the returns to experience, firms may need to deliberately manufacture the judgement and expertise that the technology no longer requires workers to practise.

5.3 Design strategic friction

The characteristic failure of this production system is uncritical acceptance, and its distinguishing feature is that it is invisible in ordinary operating data. It arrives as work delivered on time. Ford could count empty stations; a firm cannot count judgements not made, which means friction has to be built deliberately into the workflow rather than expected to emerge from it.

Concretely: require a stated position before a model is consulted on decisions above a defined materiality threshold, so that the model checks the human rather than replacing the reasoning, and place error warnings where the decision is actually made rather than in the interface periphery, where habituation erodes them. The experimental evidence supports the second intervention for deliberation while showing it does not by itself improve accuracy, which is why friction is a complement to the measurement regime in 5.1 rather than a substitute for it.

5.4 Build the moat on a public commons

This lesson concerns competitive strategy rather than labour, and it is placed here because the two turn out to be the same problem. Cheap inference confers no relative advantage: it is supplied to every competitor on the same day, and even the gap between closed-weight and freely available models has narrowed to 1.7 percent on benchmark tasks.[66]

Advantage therefore has to be built from what a vendor cannot supply, which is proprietary data, bespoke workflows, and internal feedback loops. The connection to 5.1 is the practical payoff of this section. Those assets are not produced by a strategy document; they are produced by people correcting model output, case by case, in a form the firm retains. The edit logs and rejection rationales proposed in 5.1 as a measurement mechanism are simultaneously the raw material of the moat described here. A firm that treats verification purely as a quality-control cost is paying for the asset and then discarding it. The single instruction that follows from both lessons is to capture verification work as structured data by default, because the same activity that keeps quality up is the only thing a competitor cannot buy from the same supplier.

[66] Maslej et al., AI Index Report 2025, chap. 2.

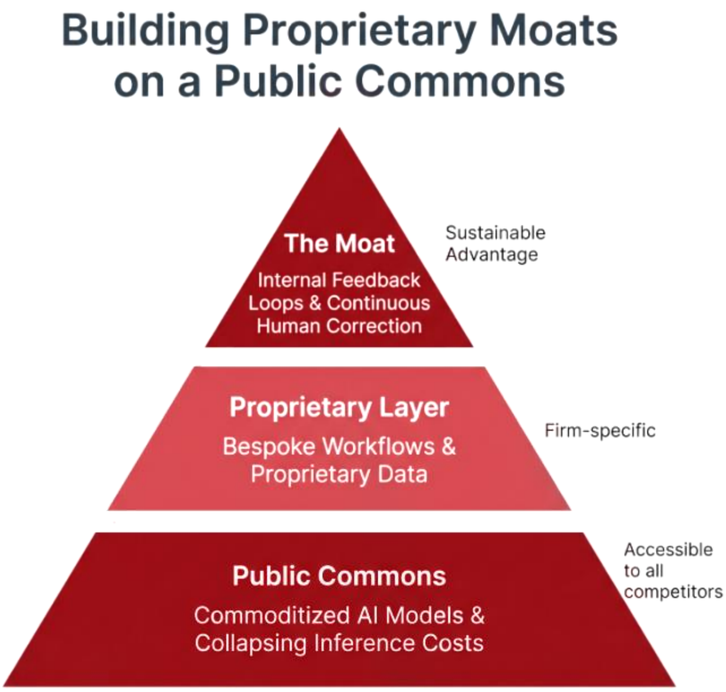

Figure 7. Building proprietary moats on a public AI commons. The durable layer sits above commoditised model access: bespoke workflows, proprietary data and internal feedback loops generated through human correction.

5.5 Move before the settlement is imposed

Ford acted unilaterally. Nobody struck, legislated, or bargained him into the five-dollar day, and the decade of price-for-quality advantage that followed reflected a lead his competitors could not close quickly.[67] Two things have changed. Settlements are already arriving in the AI economy through negotiation and regulation rather than design, which means the initiative is being lost by default. And the clock runs faster: Ford's advantage held for roughly a decade, whereas the open-to-closed model gap closed in little over a year, so any advantage built now must compound at a rate Ford never needed to achieve.

The practical implication is a sequencing one. The measurement regime in 5.1 takes time to produce a usable baseline, and the pipeline protections in 5.2 take longer still to produce their first cohort. A firm that waits for the settlement question to become urgent will be designing under conditions where the terms are set by someone else.

5.6 Counter-lesson: the surplus may not be the firm's to share

The strongest objection to this report's argument is not about implementation. It is that the Fordist analogy may fail at the point where it matters most. Ford could share his surplus because he owned it: the cost decline came from his own cumulative volume and no competitor benefited from it. The AI cost decline is supplied, arrives industry-wide simultaneously, and where competition is effective it passes through to prices rather than accumulating inside any firm. If there is no retained surplus, there is nothing to share, and the five-dollar day has no modern equivalent to design.

The objection has real force and should not be dismissed. Two things limit it. The first is that the Fordist case does not support the naive version of the argument anyway: reduced turnover and training costs offset only a small fraction of the five-dollar day's cost, and part of the improvement may be attributable to a contemporaneous downturn in the Detroit labour market rather than to the wage.[68] Ford did not pay because the saving covered the wage. He paid because the technology could not be run at full intensity by a workforce turning over at 370 percent a year, and the wage bought the acceptance that made intensification possible. The equivalent claim now is not that verification pays for itself in efficiency terms, but that a firm which cannot retain judgement cannot run its systems at the intensity they permit.

[67] Raff, "Making Cars and Making Money," 721–53.

[68] Raff and Summers, "Efficiency Wages," 29–31.

The second limit is the answer developed in 5.4. A surplus that arrives from a vendor is indeed competed away, which is precisely why the durable version has to be built from proprietary workflows, data, and feedback loops. Those are firm-specific and are not competed away on the same schedule. The counter-lesson therefore does not defeat the argument, but it does sharpen it considerably: a firm with no proprietary layer has no surplus to share and no settlement to make, and its first task is not designing a settlement but creating something to settle over.

5.7 Counter-lesson: the surveillance trap

Ford tied his money to the behaviours his system required, which was sound, and then extended the condition into workers' homes through the Sociological Department, which was not.[69] The modern form of that overreach is close enough to be worth naming precisely: monitoring how employees think with these systems rather than what they produce with them, through prompt logging, keystroke and edit-time tracking, or attention measurement.

The temptation is structural rather than incidental. Because verification quality is hard to observe directly, conduct monitoring offers a proxy that is easy to collect, and a firm that cannot measure catches will be drawn toward measuring behaviour instead. That is the failure mode 5.1 is designed to avoid, and the distinction it rests on is workable in practice: an edit log records what was changed in the work, whereas prompt surveillance records how the person was thinking. The first conditions pay on output and is defensible; the second conditions pay on conduct and, on the Fordist evidence, produces the resentment it was meant to prevent.

5.8 What a settlement would look like

Taken together, the lessons describe something reasonably specific. A firm that has made a settlement in the sense this report means pays a revocable premium for demonstrated verification rather than for assisted throughput; measures catch rates against a seeded baseline instead of trusting sign-off; budgets unassisted work as a training cost and allocates review for exposure; captures the resulting corrections as proprietary data; and does all of this before the terms are set by a strike, a regulator, or a competitor. No firm currently does all five.

The question is not whether such a settlement will be made. Instability of this kind resolves eventually, in one direction or another. The question is whether a firm makes it while the initiative is still its own, which is the one advantage Ford had in 1914 that no firm is guaranteed in 2026.

[69] Raff and Summers, "Efficiency Wages," 16–18.

References

Abernathy, William J., and Kenneth Wayne. “Limits of the Learning Curve.” Harvard Business Review 52, no. 5 (1974): 109–19.

Boston Consulting Group. Perspectives on Experience. Boston: The Boston Consulting Group, 1972.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. “Generative AI at Work.” Quarterly Journal of Economics 140, no. 2 (2025): 889–942.

Cui, Zheyuan (Kevin), Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. “The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers.” Management Science (2025). doi 10.1287/mnsc.2025.00535.

Eli, Shari, Joshua K. Hausman, and Paul W. Rhode. “The Model T.” Journal of Economic History 85, no. 1 (2025): 110–51.

Ford, Henry, and Samuel Crowther. My Life and Work. New York, 1922.

Fregin, Marie-Christine, Danique Eijkenboom, Pelin Özgül, Mantej Pardesi, and Nicholas Rounding. “Artificial Intelligence in Firms: A Synthesis of Micro-Level Evidence on Productivity, Work, and Skills.” ROA Policy Brief No. 2026/2E. Maastricht University, 2026.

Henderson, Bruce D. “The Experience Curve.” BCG Perspectives. Boston: The Boston Consulting Group, 1968.

Maslej, Nestor, et al. Artificial Intelligence Index Report 2025. Stanford Institute for Human-Centered Artificial Intelligence, 2025.

McCalley, Bruce W. Model T Ford: The Car That Changed the World. Iola, WI: Krause Publications, 1994.

Nevins, Allan. Ford: The Times, the Man, the Company. New York: Charles Scribner’s Sons, 1954. Niederhoffer,

Kate, Gabriella Rosen Kellerman, Angela Lee, Alex Liebscher, Kristina Rapuano, and

Jeffrey T. Hancock. “AI-Generated ‘Workslop’ Is Destroying Productivity.” Harvard Business

Review, 22 September 2025.

Raff, Daniel M. G. “Buying the Peace: Wage Determination Theory, Mass Production, and the Five-Dollar Day at Ford.” Revised MIT PhD thesis, 1997.

Raff, Daniel M. G. “Making Cars and Making Money in the Interwar Automobile Industry.” Business History Review 65, no. 4 (1991): 721–53.

Raff, Daniel M. G., and Lawrence H. Summers. “Did Henry Ford Pay Efficiency Wages?” NBER Working Paper No. 2101, 1986.

Schneider, Laura, and Maria Kharlamova. “Mind over Machine: Navigating Human Metacognition When Using Generative AI.” Short paper, Proceedings of the Forty-Sixth International Conference on Information Systems, Nashville, 2025.

Slichter, Sumner H. The Turnover of Factory Labor. New York: D. Appleton, 1921.