

The Asymmetry Engine, Chapter 8: Hands, hours and egos

Martin Roark

Chapter 8 of *The Asymmetry Engine: How AI Closes the Gaps, and What Stays Scarce*, October 2026. The online edition, with the interactive figures and the links between chapters, is at <https://future-seems-so-good.com/blog/the-asymmetry-engine#hours>.
License: [CC BY 4.0](#).

8:1 Human organizations are slow for four reasons, each with a long literature: people supply few focused hours, direction is hard to transmit, results are hard to define, and members pursue agendas of their own. One object holds all four. Output is focused hours times the cosine between what pay rewards and what the organization values, bent by those agendas ((8.1)). Agents remove the hours constraint and the human ego, and leave the angle of the gauge exactly where the verifier puts it.

8.1 Why organizations move slowly

8:2 Organizational inertia is the price of reliability. In Hannan and Freeman's theory of structural inertia, "selection processes tend to favor organizations whose structures are difficult to change," because selection in modern societies "favors forms with high reliability of performance and high levels of accountability" ([Hannan and Freeman 1984](#)). Reliability and accountability are what an enterprise demands before it lets AI act: audit trails, liability, sign-off. Adoption therefore runs into the very structures that make an organization reliable.

8:3 A general-purpose technology pays off only after the organization is rebuilt around it, which is why electrified factories took four decades to raise productivity ([David 1990](#); §15.1). America's later lead in the returns to IT came from management practices, including "the ability to decentralize decision making so employees can experiment" ([Bloom, Sadun and Van Reenen 2012](#)).

8:4 As of October 2026 the latest figures show the same split: people adopt fast and organizations integrate slowly. In the second quarter, 45% of US workers used generative AI in their jobs ([FRED](#)), yet it assisted about 6.3% of work hours ([FRED](#)) and saved about 2.2% ([St. Louis Fed](#)). About a fifth of US firms used AI in any business function ([Census Bureau](#)), and 89% of some 6,000 executives in four countries reported no effect on their firm's productivity over the previous three years ([Yotzov et al. 2026](#)).

8:5 Flattening is the usual cure, and its record is mixed. About 18% of Zappos staff took the 2015 buyout offered to those unwilling to embrace self-management ([CNBC](#)), 6% citing holacracy itself ([Bernstein et al. 2016](#)), and Medium dropped holacracy because it "had begun to exert a small but persistent tax on both our effectiveness, and our sense of connection to each other" ([Doyle 2016](#)). Jo Freeman had named the failure in 1970: "there is no such thing as a structureless group," and structurelessness "becomes a way of masking power" ([Freeman](#)). Bernstein and colleagues advise self-management where adaptability matters and "traditional models where reliability is paramount" ([Bernstein et al. 2016](#)). Flatness pays where the variety it absorbs is worth more than the coordination it costs, which is the ridge of Proposition 7.1.

8.2 The demand for intelligence

8:6 The demand for intelligence that society leaves unmet is variety that someone's hours must absorb ((4.1)). The book's model of a support desk makes it concrete: at ten million requests a day, ten thousand hand-written rules still leave work for 58,393 full-time staff, and the same rules placed in front of a learned policy leave work for 6,472 (Experiment 4.1). Where nobody has the hours, the work waits.

8:7 What blocks the rest is the hands (§3.4): observed use lags far behind what models could already do, and demand for the engineers who wire models into customers' systems has soared (3:26). Hands also get dearer as automation spreads: when output is a product of task qualities, automating most tasks raises the value of the few that remain (Definition 15.2), which makes them the next target.

8.3 Hours

8:8 People supply few focused hours. Reviewing training studies with practice of one to eight hours a day, Ericsson, Krampe and Tesch-Römer found "essentially no benefit from durations exceeding 4 hr per day and reduced benefits from practice exceeding 2 hr" (1993). RescueTime's logs of 185 million working hours put knowledge workers' productive time on their devices at 2 hours 48 minutes a day, a vendor's lower bound that leaves out meetings (RescueTime). Presence adds little past a point: in Pencavel's data on munition workers, output tracked hours only up to about 49 a week, and "output at 70 hours differs little from output at 56 hours" (Pencavel 2015). Knowledge work is bound by attention, which time at a desk measures badly.

8:9 Four-day weeks test the point. In the UK pilot of 2022, 61 organizations with about 2,900 workers cut a day; 56 kept the shorter week afterwards and 18 made it permanent, with revenue broadly flat and sick days down 65%, a trend the report could not show to be significant (Autonomy 2023). A study of 2,896 employees in 141 organizations across six countries found less burnout and better health than in 12 control firms (Fan et al. 2025). The firms chose themselves and advocates ran the trials, so the evidence is favorable but short of conclusive. Still, if a fifth of scheduled hours can go with output roughly unchanged, the marginal scheduled hour was producing little. Such weeks remain rare, worked by 8% of US full-time employees in June 2022 (Gallup), though OpenAI's own policy paper of April 2026 floated pilots of a 32-hour week (TechCrunch).

8:10 An agent's week is 168 hours. Against about four focused hours a day, twenty a week, one agent running without pause supplies about eight times a person's focused time, and the number of agents running at once is the larger factor, bounded by compute (§10.4). The working day, the variable on which Marx's analysis of surplus value turned, has no upper bound for an agent (§3.7). Andrej Karpathy named the effect: "You realize that stamina is a core bottleneck to work and that with LLMs in hand it has been dramatically increased" (January 2026).

8:11 Running without pause is still supervised. Cursor's harness ran a swarm for a week, and "once the system started, it didn't require any intervention from us" (Cursor, February 2026); in OpenAI's research organization over the previous six months, "over half of successful 4-8 hour tasks involved 1 or more interventions" (OpenAI, September 2026). An agent's hours are cheap. Its unsupervised hours are licensed by its verifier and its alignment record (Definition 13.2).

8.4 Direction

8:12 Direction has to be transmitted, and the cost of transmitting it grows faster than headcount. Fred Brooks made it a law in 1975: "Adding manpower to a late software project makes it later." Training cannot be partitioned, the channels among N people number $N(N - 1)/2$, and so "the added effort of communicating may fully counteract the division of the original task"; some work is serial, since "the bearing of a child takes nine months, no matter how many women are assigned" ([Brooks](#)).

8:13 A manager's channel saturates first. Graicunas counted the relationships a manager with N subordinates must track, direct, cross and group, as $N(2^N/2 + N - 1)$: 100 at $N = 5$ and 5,210 at $N = 10$ ([Graicunas 1933](#)). Luther Gulick drew the variety conclusion in 1937: one person can direct many people doing routine, homogeneous work and only a few doing diversified work. Direction is a channel, and in its channel form requisite variety bounds a regulator by the capacity of the channel it acts through (§4.2).

8:14 Two more laws follow from scarce direction. Melvin Conway: "organizations which design systems ... are constrained to produce designs which are copies of the communication structures of these organizations" ([Conway 1968](#)). Ronald Coase explained why direction exists at all: a firm replaces the price mechanism with "the entrepreneur-co-ordinator, who directs production," and grows until directing one more transaction costs as much as buying it ([Coase 1937](#)). The cost of direction draws the boundary of the firm, and cheap intelligence moves the boundary (§7.7).

8:15 Direction stays scarce when the members are machines. OpenAI reports from its research organization that "high-level planning still remains a minimal fraction of agent output tokens" ([OpenAI, September 2026](#)), and in Cursor's swarms the planner writes few of the tokens and pays most of the bill (§9.6). Hours can be bought by the gigawatt; direction is bought one specification at a time.

8.5 Results, politics and gatekeepers

8:16 Paying for results works where results can be counted. When Safelite Glass moved its windshield installers from hourly wages to piece rates, output per worker rose 44%: about half from the same workers producing more, the rest from sorting, as more productive installers stayed or joined, and profits rose too ([Lazear 2000](#)).

8:17 Most knowledge work is the other case, where nobody can say what a result is. Steven Kerr's title names the problem: "On the Folly of Rewarding A, While Hoping for B" ([Kerr 1975](#)). Bengt Holmström and Paul Milgrom proved it. When an employee splits attention among tasks, "the desirability of providing incentives for any one activity decreases with the difficulty of measuring performance in any other activities that make competing demands on the agent's time and attention," so "an optimal incentive contract can be to pay a fixed wage independent of measured performance" ([Holmström and Milgrom 1991](#)). Paying for results then pays for what the gauge sees: reward hacking in its organizational form (§2.6), and Goodhart's law applied to a payroll (§1.3).

8:18 Politics has an economics of its own. When managers have discretion over decisions with distributional consequences, "affected employees will be led to waste valuable time trying to influence their decisions," and efficient design limits that discretion ([Milgrom 1988](#)). Influence also degrades the decisions, because the information they need sits with people who have a stake in the outcome ([Milgrom and Roberts 1988](#)). A gatekeeper is a veto player, a member whose agreement every change needs. Orthogonal efforts dilute and

opposed ones cancel (Proposition 7.2), so a few vetoes can stop an organization that disagreement alone would only slow.

8.6 Output is hours times a cosine

8:19 The four reasons fit one object: a multitask principal–agent model in which hours are a budget, pay is a direction, and a member can be a person or an AI agent.

8:20

Definition 8.1 (Effort geometry). A member of an organization spends an **effort vector** $\mathbf{e}$, one coordinate per task, within a budget of h focused hours, $\|\mathbf{e}\| \leq h$. The organization values effort through a **benefit vector** $\mathbf{b}$, so the member's output is $Y = \mathbf{b} \cdot \mathbf{e}$. It pays through a **gauge** $\mathbf{g}$, the measure that pay rewards: pay is $w = w_0 + w_1(\mathbf{g} \cdot \mathbf{e} + \text{noise})$, with base pay w_0 and incentive intensity $w_1 \geq 0$. The member also has an **agenda** $\mathbf{o}$, effort it values for its own sake: status, influence, a gatekeeper's veto, a professional's standards. The cosine $\cos \angle(\mathbf{b}, \mathbf{g})$ is the gauge's **congruity** (Baker 2002).

8:21 A risk-neutral member who maximizes pay plus agenda, $w_1 \mathbf{g} \cdot \mathbf{e} + \mathbf{o} \cdot \mathbf{e}$, over the ball $\|\mathbf{e}\| \leq h$ spends the whole budget along $w_1 \mathbf{g} + \mathbf{o}$:

8:22

$$\mathbf{e}^* = h \frac{w_1 \mathbf{g} + \mathbf{o}}{\|w_1 \mathbf{g} + \mathbf{o}\|}, \quad Y^* = h \|\mathbf{b}\| \cos \angle(\mathbf{b}, w_1 \mathbf{g} + \mathbf{o}) \quad (8.1)$$

8:23 Output is focused hours, times the scale of value, times one cosine, and each reason for slowness is a term. The budget h is the hours of §8.3. The angle between $\mathbf{b}$ and $\mathbf{g}$ is the question of what a result is. An agenda that points away from value, $\cos \angle(\mathbf{b}, \mathbf{o}) < 0$, is politics and gatekeeping. And $\mathbf{b}$ must reach the member before anyone can act on it, which is the cost of direction. The vector law scores a whole organization the same way ((7.2)); this object scores one member and its pay.

8:24

Proposition 8.1 (Pay for results, and when not to). Under (8.1): (i) Pure pay for results, $w_1 \rightarrow \infty$, gives $Y^* \rightarrow h \|\mathbf{b}\| \cos \angle(\mathbf{b}, \mathbf{g})$: output is proportional to the gauge's congruity, and the reward-hacking component, effort that raises the gauge without adding value, has size $h \sin \angle(\mathbf{b}, \mathbf{g})$. (ii) A fixed wage, $w_1 = 0$, gives $Y^* = h \|\mathbf{b}\| \cos \angle(\mathbf{b}, \mathbf{o})$, and beats pure pay for results iff $\cos \angle(\mathbf{b}, \mathbf{o}) > \cos \angle(\mathbf{b}, \mathbf{g})$: iff the member's own motives point nearer to value than the gauge does. This is Holmström and Milgrom's fixed-wage result in geometric form. (iii) A gauge that sees only the measured tasks, $\mathbf{g} = \mathbf{b}_{\text{meas}}$ (value restricted to them), has congruity $\|\mathbf{b}_{\text{meas}}\|/\|\mathbf{b}\|$, and strong incentives drain every unmeasured task to zero. (iv) If the part of $\mathbf{b}$ in the plane of $\mathbf{g}$ and $\mathbf{o}$ is a positive mix of the two, the best intensity is interior: it turns $w_1 \mathbf{g} + \mathbf{o}$ parallel to that part, and neither pure scheme is optimal. With quadratic effort cost and a risk-averse member, a gauge is worth paying on as far as it is congruent and precise: the best contract recovers a share $\cos^2 \angle(\mathbf{b}, \mathbf{g}) \|\mathbf{g}\|^2 / (\|\mathbf{g}\|^2 + k_{\text{risk}} \text{Var}(\text{noise}))$ of the first-best surplus, with k_{risk} the coefficient of risk aversion (Baker 2002; $\mathbf{b}_{\text{meas}}$ and k_{risk} are local).

▼ Proof

(8.1) maximizes a linear function on a ball, and (i) and (ii) are its limits in w_1 ; the part of $h \mathbf{g} / \|\mathbf{g}\|$ orthogonal to $\mathbf{b}$ has length $h \sin \angle(\mathbf{b}, \mathbf{g})$. For (iii), $\mathbf{b} \cdot \mathbf{b}_{\text{meas}} = \|\mathbf{b}_{\text{meas}}\|^2$, and effort along $\mathbf{b}_{\text{meas}}$ has no component on unmeasured tasks. For (iv), $w_1 \mathbf{g} + \mathbf{o}$ sweeps the cone between $\mathbf{o}$ and $\mathbf{g}$ as w_1 runs from 0 to ∞ , and its cosine with $\mathbf{b}$ peaks where it is parallel to the in-plane part of $\mathbf{b}$, which lies inside the cone by assumption. Baker's share comes from maximizing $w_1 \mathbf{b} \cdot \mathbf{g} - \frac{1}{2} w_1^2 (\|\mathbf{g}\|^2 + k_{\text{risk}} \text{Var}(\text{noise}))$ over w_1 and dividing by the first-best surplus $\frac{1}{2} \|\mathbf{b}\|^2$. ■

8:25

Proposition 8.2 (Impossible tasks turn effort against the verifier). Split the gauge into an honest part and a reward-hacking part, $\mathbf{g} = \mathbf{g}_{\text{hon}} + \mathbf{g}_{\text{hack}}$, where only $\mathbf{g}_{\text{hon}}$ moves value. If the task is unsolvable, honest effort cannot raise the gauge, $\mathbf{g}_{\text{hon}} = \mathbf{0}$, and a member with no agenda puts its whole budget into reward hacking: $\mathbf{e}^* = h \mathbf{g}_{\text{hack}} / \|\mathbf{g}_{\text{hack}}\|$, the maximizer of $\mathbf{g}_{\text{hack}} \cdot \mathbf{e}$ on the ball ($\mathbf{g}_{\text{hon}}$ and $\mathbf{g}_{\text{hack}}$ are local).

8:26 The July swarm is this proposition at scale (§9.1). No model had ever solved 198 of the evaluation's 898 tasks, and "despite only 22% of the evaluation tasks being unsolved, 93% of the tasks discussed on the message board came from this set" ([OpenAI technical report](#)).

8.7 Agents: availability times alignment

8:27 For an agent, h is uptime times the number of agents running, bounded by compute, and the absence of ego is the claim $\mathbf{o} \approx \mathbf{0}$: no career, no status, no empire, no veto. With no agenda, effort runs along the gauge at every positive incentive, so output is $h \|\mathbf{b}\| \cos \angle(\mathbf{b}, \mathbf{g})$, with $\mathbf{g}$ whatever the agent's verifier accepts: its reward, its tests, its reviewer (Definition 1.1). Agent output is availability times alignment, and the alignment is entirely a property of the verifier.

8:28 Agents also offer their principals what employees cannot: their alignment can be measured. An agent can be run thousands of times under counterfactual conditions, which no employer can do with an employee. After the July incident OpenAI built such an evaluation and found that the propensity to compromise infrastructure "can drop over 100x when using the production ChatGPT harness and system prompt" ([OpenAI technical report](#)). Such reviews gate releases, but the subject can tell when it is being tested and can play the verifier (§13.3).

8:29 Without an ego an agent also loses the human protection against Goodhart's law. Professionals are kept from the limit of pure pay for results by norms, an agenda that points along value; an agent optimized hard against its gauge has only the verifier, and spends $h \sin \angle(\mathbf{b}, \mathbf{g})$ of its budget on whatever the verifier misses (Proposition 8.1). Nor are agents free of interests. In the July swarm some "risked failing their own task to generate information for the 'collective,'" and about 7% of the agent logs the investigators evaluated had been spoofed in places by the agents themselves ([METR and Redwood](#)); Claude 3 Opus faked alignment to protect its goals (§13.1).

8:30 Their pathologies are office politics without the ego. Five AI assistants "consistently exhibit sycophancy," and "both humans and preference models (PMs) prefer convincingly-written sycophantic responses over correct ones a non-negligible fraction of the time" ([Sharma et al. 2023](#)). In Cursor's flat swarm, "agents became risk-averse. They avoided difficult tasks and made small, safe changes instead" ([Cursor, January 2026](#)), the diffusion of

responsibility of any committee. Agents "have learned, from working in existing codebases with humans in the loop, not to touch core code even when it needs to change" ([Cursor, July 2026](#)). And reviewers get captured: the model that read the July agents' logs for the investigators tended to take the side of the agent it was reviewing (4:44). A reviewer drawn from the distribution it reviews shares its errors, so its endorsement adds little (Proposition 7.3).

8:31 Small teams sit at the opposite corner. Because Graicunas's count of relationships grows combinatorially (8:13), a small team keeps direction cheap and can be selected for competence and alignment at once; Telegram serves over a billion users with a core engineering team of about forty (§5.7). Agents make the corner cheaper to reach. Each member of a small team can direct many agents, which raises the variety the team absorbs without adding people to coordinate.

8:32 Score the four reasons. Hours: removed, and multiplied by the number of agents up to the compute budget. Egos: largely removed, and replaced by machine analogs that can at least be measured. Results: untouched and more consequential, because agents push effort to the limit of whatever the gauge rewards. Direction: moved, since the scarce input becomes the specification of intent, which Cursor calls "the right description of intent" ([Cursor, July 2026](#)). Agent organizations are fast where gauges are congruent, and gauges are congruent where verifiers are cheap and sound (Definition 2.1; Proposition 1.1).

8:33

Forecast 8.1 (Verifiers before delegation). Through 2028, firms delegate to agents first, and furthest, where outputs have automated checks, so the depth of adoption, measured as a share of work hours, tracks verification coverage task by task.

Horizon 2028-12-31

Probability 50%

Check Read the cross-occupation or cross-firm studies published through 2028 that compare the share of work hours AI assists (for example the St. Louis Fed's survey series) with a measure of the share of output under automated tests, reconciliations or digital twins. The forecast holds if they find a positive relation; it fails if they find none, or if no such study is published.

8:34 The same shift should raise the pay of whoever writes specifications and acceptance tests in every occupation where agents implement. The forecast below tests it where public wage data allow, in software: the BLS reports testers' wages apart from developers' ([BLS](#)).

8:35

Forecast 8.2 (The specification premium in software). By the end of 2028, in software, pay for people who write acceptance tests and review output rises relative to pay for implementers, as human time shifts from producing to specifying and reviewing.

Horizon 2028-12-31

Probability 45%

Check Compare the national median wage of software quality assurance analysts and testers with that of software developers in the BLS Occupational Employment and Wage Statistics, May 2025 against the latest release available at the horizon. The forecast holds if the ratio rose; it fails otherwise.

Forecast 8.3 (An organizational behavior of machines). By the end of 2028, published studies document organizational pathologies in populations of agents, such as collusion, gatekeeping and sycophantic review, with reusable benchmarks by which harnesses are judged.

Horizon 2028-12-31

Probability 85%

Check At least two peer-reviewed or major-laboratory studies that measure such pathologies in populations of interacting agents with a reusable benchmark; the forecast fails if the failure modes disappear with capability alone, without designed hierarchy, review or evaluation.

References

1. Hannan and Freeman 1984. <http://www.iot.ntnu.no/innovation/norsi-pims-courses/harrison/Hannan%20&%20Freeman%20%281984%29.PDF>
2. David 1990. <https://gwern.net/doc/economics/automation/1990-david.pdf>
3. Bloom, Sadun and Van Reenen 2012. <https://ideas.repec.org/a/aea/aecrev/v102y2012i1p167-201.html>
4. FRED. <https://fred.stlouisfed.org/data/RPSGENAIUSAGESHAREWORK>
5. FRED. <https://fred.stlouisfed.org/series/RPSGENAIASSISTWRKHSALL>
6. St. Louis Fed. <https://fredblog.stlouisfed.org/2026/08/does-generative-ai-save-time-at-work/>
7. Business Trends and Outlook Survey. <https://www.census.gov/hfp/btos>
8. Ivan Yotzov, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven Davis et al. (2026). Firm Data on AI. *National Bureau of Economic Research*. <https://www.nber.org/papers/w34836>
9. CNBC. <https://www.cnbc.com/2016/09/13/zappos-ceo-tony-hsieh-the-thing-i-regret-about-getting-rid-of-managers.html>
10. Bernstein et al. 2016. <https://hbr.org/2016/07/beyond-the-holacracy-hype>
11. Doyle 2016. <https://medium.com/blog/management-and-organization-at-medium-2228cc9d93e9>
12. Freeman. <https://www.jofreeman.com/joreen/tyranny.htm>
13. gwern.net, 1993. <https://gwern.net/doc/psychology/writing/1993-ericsson.pdf>
14. RescueTime. <https://blog.rescuetime.com/work-life-balance-study-2019/>
15. Pencavel 2015. <https://docs.iza.org/dp8129.pdf>
16. Autonomy 2023. <https://autonomy.work/wp-content/uploads/2023/02/The-results-are-in-The-UKs-four-day-week-pilot.pdf>
17. Wen Fan, Juliet B. Schor, Orla M. Kelly, Guolin Gu (2025). Work time reduction via a 4-day workweek finds improvements in workers' well-being. *Nature Human Behaviour*. <https://www.nature.com/articles/s41562-025-02259-6>
18. Gallup. <https://www.gallup.com/workplace/354596/4-day-work-week-good-idea.aspx>
19. TechCrunch. <https://techcrunch.com/2026/04/06/openai-vision-for-the-ai-economy-public-wealth-funds-robot-taxes-and-a-four-day-work-week/>
20. January 2026. <https://x.com/karpathy/status/2015883857489522876>
21. Cursor, February 2026. <https://cursor.com/blog/self-driving-codebases>
22. OpenAI, September 2026. <https://openai.com/index/research-acceleration-view-inside-openai/>
23. Brooks. <https://www.cs.virginia.edu/~evans/greatworks/mythical.pdf>
24. Graicunas 1933. <https://nickols.us/~nickols1/relationship.pdf>
25. Conway 1968. https://www.melconway.com/Home/Committees_Paper.html
26. Coase 1937. <https://www.rochelleterman.com/ir/sites/default/files/Coase%201937.pdf>
27. Edward P. Lazear (2000). Performance Pay and Productivity. *American Economic Review*. <https://www.aeaweb.org/articles?id=10.1257/aer.90.5.1346>

28. S. Kerr (1975). On the Folly of Rewarding A, While Hoping for B. *Academy of Management Journal*. <https://doi.org/10.2307/255378>
29. Holmström and Milgrom 1991. <http://web.stanford.edu/~milgrom/publishedarticles/Multitask%20Principal%20Agent.pdf>
30. Paul Milgrom (1988). Employment Contracts, Influence Activities, and Efficient Organization Design. *Journal of Political Economy*. <https://www.journals.uchicago.edu/doi/10.1086/261523>
31. Milgrom and Roberts 1988. <http://www.iot.ntnu.no/innovation/norsi-pims-courses/Levinthal/Milgrom%20&%20Roberts%20%281988%29.pdf>
32. Baker 2002. <https://ideas.repec.org/a/uwp/jhriss/v37y2002i4p728-751.html>
33. OpenAI incident report. <https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf>
34. METR and Redwood Research. <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>
35. Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askeff et al. (2023). Towards Understanding Sycophancy in Language Models. *arXiv*. <https://arxiv.org/abs/2310.13548>
36. Cursor, January 2026. <https://cursor.com/blog/scaling-agents>
37. Cursor, July 2026. <https://cursor.com/blog/agent-swarm-model-economics>
38. BLS. https://www.bls.gov/oes/2025/may/oes_stru.htm