

CENTER FOR MOTIVATION RESEARCH

AI as a Time Trap

A systems-thinking review of how artificial intelligence can consume the time it promises to save

Author: John Angheli

Contents

Executive summary

1. Introduction

2. The two-income trap as a systems template

3. The AI time-trap loop

4. Mechanisms in ranked order

5. Evidence and quantified patterns

6. How to avoid the trap

7. Open questions and limitations

8. Conclusion

References

Executive summary

The **AI time trap** is the pattern in which AI lowers the time cost of an individual task, but the surrounding system adapts so that the saved minutes are swallowed by higher expectations, extra review, more inbound work, and new coordination burdens. In the same way that the two-income trap turns a seeming household surplus into a new baseline of fixed costs, AI can turn a labour-saving tool into a new minimum standard of responsiveness, output, and availability (Warren and Tyagi, 2003).

The trap is not that AI makes people slower in general. The evidence is more exacting than that. In bounded tasks, AI can create genuine gains. Noy and Zhang's controlled writing study found completion times fell by 40% and output quality rose by 18% (Noy and Zhang, 2023). Brynjolfsson, Li and Raymond found 14% higher productivity on average in customer support, with larger gains for novice agents (Brynjolfsson et al., 2023). Dillon and colleagues found regular users in a cross-industry field experiment spent substantially less time on email (Dillon et al., 2025).

Yet broader workplace evidence shows why these gains often fail to become free time. In Humlum and Vestergaard's Danish worker panel, average reported time savings among users were only 2.8% of work hours, and 80% of those gains were reallocated into other job tasks (Humlum and Vestergaard, 2025).

Workday's employer survey found nearly 40% of AI time savings lost to rework and only a minority of employees reporting consistently clear net benefits (Workday, 2026). Upwork's workforce survey found many workers saying AI had increased workload in at least one way (Upwork Research Institute, 2024). METR's randomized study of experienced open-source developers found that early-2025 AI tools made them 19% slower on measured tasks (Becker et al., 2025).

The most important mechanisms, in order, are the throughput ratchet, the verification tax, workflow mismatch, new meta-work and fixed costs, task misfit along AI's jagged frontier, loss of backup human capability, AI debt, and status or competitive pressure. The best solutions are correspondingly systemic: measure net end-to-end time rather than draft speed; redesign workflows instead of only adding tools; use AI to reduce inbound load, not merely accelerate outbound load; restrict AI first to bounded tasks; train users to verify selectively and critically; preserve human judgment and skill; and convert some savings into real slack rather than into higher quotas (Calvino et al., 2025; Workday, 2026).

1. Introduction

Artificial intelligence is usually sold under the grammar of liberation. It will write the email, summarize the meeting, draft the proposal, clean the spreadsheet, prepare the first version of the code, and reduce the cognitive friction of daily work. This promise is not empty. There are now strong studies showing that AI can improve performance on specific tasks, especially when the task is bounded, the quality standard is clear, and the user can judge the result (Noy and Zhang, 2023; Brynjolfsson et al., 2023).

The problem is that time is not only an individual possession. Time is also a system property. A person may save ten minutes on a draft, but the team may use that ten minutes to raise the expected number of drafts. A salesperson may write a better email faster, but if every competitor does the same, the buyer's inbox becomes a harsher battlefield. A clinician may save documentation time, but the practice may convert that time into more visits rather than a shorter day. The saved minute is real. The question is who captures it.

This is why the analogy to the two-income trap is useful. In that argument, the second income should have created household resilience. Instead, many families used the new income to support higher fixed costs - larger mortgages, school-district competition, childcare, transport, and other non-negotiable expenses. The second income became a requirement rather than a buffer (Warren and Tyagi, 2003).

AI can follow a similar path. It begins as extra capacity. Then organizations and markets adjust. Output norms rise. Communication volume expands. Rework grows. Human skills atrophy. Tooling, training, compliance, and integration become recurring obligations. The assistant becomes part of the baseline.

The useful question, therefore, is not simply: *Does AI save time on the task?* The better question is: *After the whole system adapts, does the user or organization have more discretionary time, clearer attention, better judgment, and lower coordination burden?*

2. The two-income trap as a systems template

The two-income trap is powerful because it exposes a counterintuitive systems pattern: an apparent increase in capacity can make the actor more fragile if the

surrounding system adapts to absorb the gain. A second income is not harmful in itself. It becomes harmful when the household's fixed-cost structure rises until both incomes are required merely to stand still.

The same structure appears in technology adoption. A tool that reduces unit cost often increases total use. This is the familiar rebound logic often associated with Jevons's analysis of coal efficiency: making a resource cheaper to use can increase aggregate consumption of that resource rather than reduce it (Jevons, 1865). With AI, the scarce resource is not only money or energy. It is attention, judgment, verification capacity, and organizational bandwidth.

The table below translates the two-income pattern into the AI time-trap analogue.

Table 1. Systems parallels between the two-income trap and the AI time trap

Two-income trap pattern	AI analogue	What it looks like in practice
Bidding / ratchet	Throughput ratchet	Cheaper drafting raises expectations for faster replies and more deliverables. Saved time becomes expanded work (Humlum and Vestergaard, 2025).
Loss of backup	Loss of unaided human capacity	Routine judgment is offloaded; critical thinking, problem framing, and exception handling receive less practice (Lee et al., 2025).
Fixed costs	AI fixed-cost layer	Review queues, compliance, prompt libraries, governance, tool subscriptions, and integration become recurring obligations (Workday, 2026).
Verification burden	Rework tax	Checking, correcting, rewriting, and validating can absorb much of the saved time; Workday estimates nearly 40% is lost here (Workday, 2026).
Competition escalation	Communication glut and norm inflation	As messages and proposals become cheaper to produce, attention becomes scarcer and screening burdens rise (Microsoft WorkLab, 2025).
Debt / credit analogue	Content, technical, and decision debt	AI gives instant draft liquidity, then creates future obligations to integrate, review, reconcile, and clean up (Calvino et al., 2025).
Status / guilt pressure	Mandates and fear of falling behind	AI use becomes professionally compulsory, even where the use case has not yet proven net savings (Upwork Research Institute, 2024).

3. The AI time-trap loop

The causal loop is the heart of the trap. The gain at the individual level feeds a collective ratchet that can destroy slack. This is why the problem is better understood as a systems effect than as a simple tool defect.

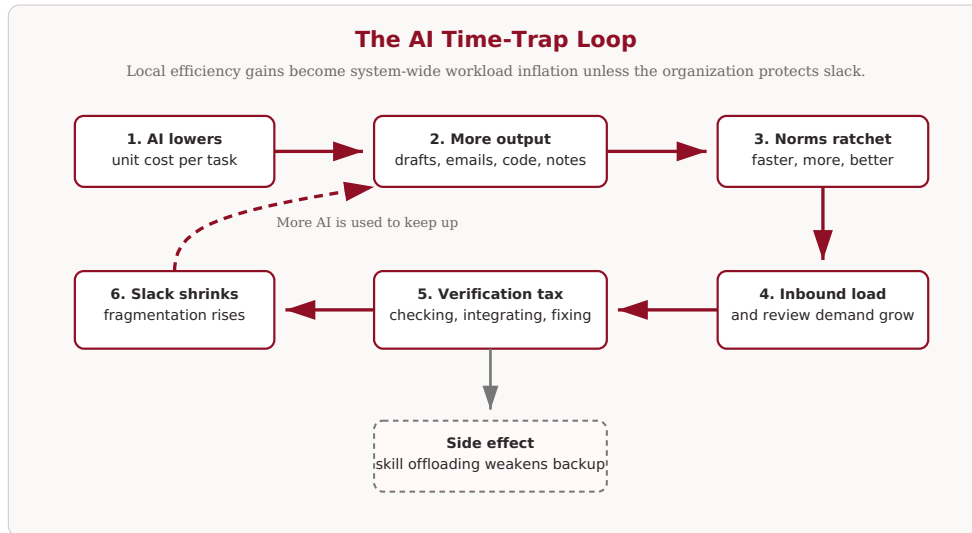


Figure 1. AI lowers the unit cost of work, but the broader system can convert that gain into higher output norms, more inbound work, and heavier verification demand.

The loop begins with a real efficiency gain. The first draft is faster. The summary appears instantly. The code suggestion arrives before the developer has fully searched the repository. But lower unit cost changes behaviour. People produce more. Managers expect more. Recipients receive more. Then the system must spend time evaluating, prioritizing, correcting, and coordinating all that extra production.

This is why the email example is analytically useful even though direct causal evidence on AI-generated email volume remains limited. The background conditions are already saturated. Microsoft WorkLab reports that the average worker receives high volumes of email and chat, with frequent interruptions during core work hours (Microsoft WorkLab, 2025). In that environment, cheaper outbound communication can easily worsen the scarcity of attention unless the organization explicitly reduces inbound load.

4. Mechanisms in ranked order

The mechanisms below are ranked by structural importance rather than by moral severity. The first three are the most decisive because they determine whether task speed survives as usable time.

① Throughput ratchet

AI makes output cheaper, so the definition of normal output rises. Humlum and Vestergaard found that most reported time savings were reallocated into other tasks rather than captured as slack (Humlum and Vestergaard, 2025).

Metaphor: a moving walkway helps only until everyone is required to walk faster on it.

② Verification tax

AI shifts work from production to checking. The draft arrives quickly, but the human must still verify facts, tone, legal risk, technical soundness, and fit. Workday reports a large share of savings lost to rework (Workday, 2026).

Metaphor: borrowing against future attention.

③ Workflow mismatch

AI can accelerate individually controlled tasks while leaving meetings, approvals, handoffs, and queues untouched. Dillon and colleagues found email time fell, but meeting time did not significantly change (Dillon et al., 2025).

Metaphor: a turbocharger on a traffic jam.

④ New meta-work and fixed costs

Tool access creates new obligations: integration, prompts, governance, training, compliance, audit, and model maintenance. This is the AI equivalent of taking on a new fixed household expense (Humlum and Vestergaard, 2025; Workday, 2026).

5 Task misfit on the jagged frontier

AI is excellent on some tasks and poor on others. The same tool can be a bicycle on pavement or a bicycle in sand. Misfit creates false starts and cleanup (Dell'Acqua et al., 2023; Calvino et al., 2025).

6 Loss of backup human capacity

When AI handles the routine reasoning too often, the human reserve weakens. Lee and colleagues found higher confidence in GenAI associated with less critical thinking (Lee et al., 2025).

7 AI debt

AI can create drafts, code, notes, and decisions that look finished but require later reconciliation. The debt is paid in bug fixing, content cleanup, review queues, and lost trust.

8 Status and competitive pressure

Once AI becomes a signal of diligence, people use it even where the use case is weak. Upwork's survey suggests many workers feel higher expectations without adequate redesign or training (Upwork Research Institute, 2024).

Two boundary cases

Software development shows why speed can be misleading. AI may produce code quickly, but the real work includes integration, testing, security, and maintainability. METR's randomized study of experienced open-source developers found a slowdown in that expert setting (Becker et al., 2025).

Medical documentation shows the same pattern from another angle. AI scribes can reduce EHR and documentation time, yet Rotenstein and colleagues found no significant reduction in after-hours EHR time (Rotenstein et al., 2026). Local efficiency improved, but personal time did not clearly return.

5. Evidence and quantified patterns

The empirical record is easiest to interpret when separated into **task-level gains** and **system-level capture**. Task-level studies often show impressive

speed. Workplace studies ask the more difficult question: where did the saved time go?

Table 2. Evidence map: gross gain versus hidden capture

Study and setting	Gross gain	Hidden cost or capture	Net lesson
Noy and Zhang: professional writing tasks	40% faster; 18% higher quality	Editing and validation still required	Strong gains in bounded tasks, not proof of lighter schedules (Noy and Zhang, 2023).
Brynjolfsson, Li and Raymond: customer support	14% average gain; larger novice gains	Smaller gains for top performers	AI often helps less-experienced workers most (Brynjolfsson et al., 2023).
Dillon et al.: cross-industry field experiment	Less email time for regular users	Meetings unchanged	Individual speedups do not automatically fix shared bottlenecks (Dillon et al., 2025).
Humlum and Vestergaard: exposed occupations	2.8% average time savings among users	80% reallocated to other tasks; new AI workloads reported	Most savings are recaptured by work (Humlum and Vestergaard, 2025).
Workday global employer survey	Reported weekly time savings	Nearly 40% lost to rework; few consistently clear net outcomes	Draft speed frequently overstates real ROI (Workday, 2026).
Upwork workforce survey	High leadership expectations	Many workers report added workload or decreased productivity in at least one way	Poor rollout turns AI into more work (Upwork Research Institute, 2024).
METR: experienced open-source developers	No measured gain in the study context	19% slowdown	Expert workflows can be net-negative when AI does not fit the task (Becker et al., 2025).
Rotenstein et al.: AI scribes in healthcare	Reduced EHR and documentation time	No significant change in after-hours EHR time	Task savings do not guarantee reclaimed personal time (Rotenstein et al., 2026).

The following graphic replaces the confusing beta chart with a simple time-budget illustration. It takes the reported 2.8% saving among users and shows how little becomes protected slack once most of the saving is reabsorbed into other work.



Figure 2. Gross AI time savings are not the same as time wealth. Based on Humlum and Vestergaard's reported distribution of time savings (Humlum and Vestergaard, 2025).



Figure 3. The evidence base increasingly suggests that task gains are real, but their conversion into free time depends on workflow design and organizational norms.

6. How to avoid the trap

The core governance rule is simple: **measure net time, not tool time**. The unit of analysis should be end-to-end cycle time, including prompting, checking, corrections, approvals, downstream confusion, and exception handling. If the only visible improvement is a faster first draft, the organization is probably measuring the wrong thing (Workday, 2026).

Table 3. Practical safeguards for human-AI work

Design principle	What it requires	Anti-trap effect
Measure end-to-end time	Track the full cycle from task start to final accepted output, including review and rework.	Prevents first-draft speed from being mistaken for real time saved.
Redesign workflows	Change meeting norms, review thresholds, approval chains, escalation rules, and ownership boundaries.	Allows local task speed to propagate into system-level time savings (Dillon et al., 2025).
Reduce inbound load	Use AI for triage, prioritization, summarization, retrieval, and filtering before using it to send more output.	Protects attention instead of flooding the system with more content.
Start with bounded tasks	Deploy AI first where goals are clear and correctness can be cheaply checked.	Stays inside the useful side of the jagged frontier (Calvino et al., 2025).
Preserve human backup	Keep humans responsible for problem framing, criteria, evidence inspection, and final judgment in high-stakes work.	Prevents skill erosion and brittle dependence (Lee et al., 2025).
Use useful friction	Ask AI to challenge assumptions, flag uncertainty, and present counterarguments rather than only smoothing output.	Maintains critical engagement and reduces passive acceptance (Drosos et al., 2025).
Convert gains into slack	Protect some savings as fewer meetings, slower response norms, narrower queues, or explicit focus blocks.	Stops every gain from being capitalized into higher quotas.

For leaders

Leaders should treat AI as an operating-system change, not a plug-in. That means every major deployment should answer five questions: What work will disappear? What work will increase? Who verifies the output? Which communication norms will be reduced? What portion of the gain will become real slack rather than higher throughput?

For individuals

Individuals should treat AI as a lever for boring work, not a substitute for core judgment. Use it for triage, first drafts, transcription, retrieval, formatting, compression, and routine transformation. Be more cautious with strategic choices, domain analysis, final prose, novel edge cases, and any work where a

subtle mistake would be hard to detect. If you cannot cheaply verify the output, you probably should not outsource the step.

PRACTICAL RULE

AI first draft; human final model of reality. Let AI lower the cost of exploration, but do not let it own the standard of truth.

7. Open questions and limitations

The broad pattern is clear enough to guide action, but several questions remain unsettled.

First, the email-glut example is analytically strong but still under-measured. Existing communication-overload data supports the concern, yet direct causal evidence that AI-generated email has raised internal email volumes is still limited (Microsoft WorkLab, 2025).

Second, model progress matters. Some measured slowdowns from 2025 may narrow as tools improve, workflows mature, and users learn where AI is useful. The correct conclusion is not that AI wastes time, but that its time effect is conditional on task choice, human skill, verification burden, and workflow design (Calvino et al., 2025).

Third, distribution matters. Not every setting produces a trap. In some environments, AI may be captured as quality improvement, reduced stress, faster learning, or genuine slack. The critical question is whether organizations design for those outcomes or simply increase throughput.

Future research should track long-run equilibrium effects, direct measures of communication volume, randomized tests of workflow redesign, and the preservation of human skill under routine AI use.

8. Conclusion

The AI time trap is a warning against naive efficiency thinking. Artificial intelligence can absolutely save time at the task level. But task time is not the same as life time, workday time, or organizational slack. Between the individual task and the lived day stand norms, incentives, coordination burdens, expectations, verification, and competition.

The central lesson from the two-income analogy is that capacity gains must be protected. A second income becomes a trap when it is absorbed into fixed costs. AI becomes a trap when speed is absorbed into more messages, more deliverables, more review, more fragmented attention, and higher performance baselines.

The answer is not to reject AI. The answer is to govern it. Use it where its strengths are clear. Measure the whole workflow. Preserve human judgment. Reduce inbound noise. Keep some of the gain as genuine slack. Otherwise, the apparent liberation of AI will harden into a new form of time poverty.

References

Becker, J. et al. (2025) *Measuring the impact of early-2025 AI on experienced open-source developer productivity*. METR Report.

Brynjolfsson, E., Li, D. and Raymond, L.R. (2023) *Generative AI at work*. NBER Working Paper No. 31161. Cambridge, MA: National Bureau of Economic Research.

Calvino, F., Reijerink, J. and Samek, L. (2025) *The effects of generative AI on productivity, innovation and entrepreneurship*. OECD Artificial Intelligence Papers No. 39. Paris: OECD Publishing.

Dell'Acqua, F. et al. (2023) *Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality*. Harvard Business School Working Paper No. 24-013.

Dillon, E.W., Jaffe, S., Immorlica, N. and Stanton, C.T. (2025) *Shifting work patterns with generative AI*. NBER Working Paper No. 33795. Cambridge, MA: National Bureau of Economic Research.

Drosos, I. et al. (2025) 'It makes you think: Provocations help restore critical thinking to AI-assisted knowledge work', *arXiv*, arXiv:2501.17247.

Humlum, A. and Vestergaard, E. (2025) *Large language models, small labor market effects*. NBER Working Paper No. 33777. Cambridge, MA: National Bureau of Economic Research.

Jevons, W.S. (1865) *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*. London: Macmillan.

Lee, H.-P.H. et al. (2025) 'The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers', *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

Microsoft WorkLab (2025) *Breaking down the infinite workday*. Microsoft.

Noy, S. and Zhang, W. (2023) 'Experimental evidence on the productivity effects of generative artificial intelligence', *Science*, 381(6654), pp. 187-192.

Olson, K.D. et al. (2025) 'Use of ambient AI scribes to reduce administrative burden and professional burnout', *JAMA Network Open*.

Rotenstein, L.S. et al. (2026) 'Changes in clinician time expenditure and visit quantity with adoption of artificial intelligence-powered scribes: A multisite study', *JAMA*.

Upwork Research Institute (2024) *From burnout to balance: AI-enhanced work models for the future*. Upwork.

Warren, E. and Tyagi, A.W. (2003) *The Two-Income Trap: Why Middle-Class Parents Are Going Broke*. New York: Basic Books.

Workday (2026) *Beyond productivity: Measuring the real value of AI*. Workday Research Summary.
