

Hedegreen Research

The Saved Hours Doctrine

Productivity, AI and the Question of Who Captures Time

Author	Dennis Hedegreen
Publisher	Hedegreen Research
Date	May 2026
Status	Pre-paper review draft
Version	0.93
Issue	I001
Type	working-paper

Local paper build generated from the Hedegreen Research markdown source.

Contents

Abstract	1
Keywords	1
1. Introduction	1
1.1 Human dignity and the AI social question	2
2. Conceptual Model	2
2.1 Saved-hours audit protocol	3
3. Definitions	4
4. Output Definition and Metric Capture	4
5. Uneven Compressibility	5
6. Productivity Risk	5
7. Capture Framework	6
8. Case Tests	7
8.1 China: productivity as state and industrial capacity	7
8.2 Singapore: productivity as wage and skills bargain	7
8.3 United States: measurement without distribution	8
8.4 European Union: social dialogue and algorithmic management .	8
8.5 Greece: the reverse case	9
8.6 Denmark: measurement and public-sector comparator	10
9. Policy Implications	11
10. Counterarguments and Failure Modes	12
11. Limitations	13
12. Conclusion	13
References	14
Appendix Candidate: School Pilot On Saved-Time Intuitions	16

Abstract

Public debate about artificial intelligence and work often begins with job loss. This paper argues that an earlier question should be made visible: when AI increases productivity, can it take hours before it takes jobs? The paper introduces the concept of saved hours: human labour time made technically unnecessary for maintaining a defined baseline of output after a demonstrated productivity increase. Saved hours are not automatically free time. They may be captured as profit, output, lower prices, state capacity, wage progression, labour displacement, intensified work, or reduced working time. The paper develops a simple measurement model, distinguishes saved hours from time dividends, and identifies capture mechanisms including metric capture and algorithmic management. It then uses China, Singapore, the United States, the European Union, Greece and Denmark as case tests for different institutional pathways. The central claim is that AI governance should not only measure jobs lost or GDP gained. It should also measure hours saved, hours captured and hours returned.

Keywords

artificial intelligence; productivity; working time; labour policy; automation; algorithmic management; public sector; time use; four-day week; saved hours

1. Introduction

The standard labour question around artificial intelligence is whether AI will take jobs. That question matters, but it is not the first measurable question. Before a job disappears, the hours inside the job may already have changed.

AI systems can reduce the time needed to draft documents, summarise information, translate routine material, prepare reports, classify tickets, answer customer queries, process invoices, write code, search case material or complete administrative tasks. Where the same baseline output can be produced in fewer human hours, a productivity surplus has appeared before the employment contract itself has necessarily changed.

This paper calls that surplus saved hours.

The core proposition is narrow. AI does not automatically create unemployment or freedom. It may create a technical reduction in the human time required for a task, workflow or sector at a defined output baseline. The distribution of that saved time is then an institutional question.

The existence of saved hours is technical. The distribution of saved hours is political.

The purpose of this paper is to make that distribution problem visible. If AI policy measures only jobs lost, GDP gained or firm-level productivity, it may

miss the earlier capture of labour time. Where saved hours exist, they can disappear into higher output expectations, platform targets, monitoring systems, public-sector throughput, state capacity, wages, prices or profit before they are ever visible as time.

1.1 Human dignity and the AI social question

After V0.92 of this paper was made public as a Camelot preflight draft, Pope Leo XIV's **Magnifica Humanitas** placed artificial intelligence explicitly inside the living tradition of Catholic social doctrine. The encyclical should not be treated as evidence for the saved-hours model, and this paper does not depend on a theological frame. Its relevance is contextual. It shows how quickly AI has become a question not only of technical risk, but of human dignity, work, education, social power and the common good.

The overlap is narrow but important. **Magnifica Humanitas** warns against reducing the mystery of the person to data and performance, against technological power being concentrated in private actors, and against systems whose effects on employment, services, data management and daily life are imposed without transparency, participation or recourse. It also places schools and education inside the digital transition, not outside it.

The saved-hours doctrine asks the same broad social question in a more limited labour-policy form. If AI creates a measurable productivity surplus, human dignity is not protected merely by naming the tool efficient. The surplus must be made visible. A worker, student, patient, citizen or public servant should not be reduced to an output variable whose time can be silently reallocated by whoever controls the metric. In this sense, the saved-hours audit is one practical way to ask a dignity question in measurable form: when human time has been made technically unnecessary, who decides what that time becomes?

2. Conceptual Model

The model begins with a deliberately simple relationship:

$$O = H \times P$$

where O is output, H is human labour hours in a defined period and P is output per human labour hour. This is a conceptual identity rather than a sector-specific production function.

For a constant baseline of output, saved hours can be defined as:

$$S = H_0 - H_1$$

where S is saved hours, H_0 is the old number of human hours required for the baseline output and H_1 is the new number of human hours required for the same baseline after a productivity increase.

The required productivity multiplier for a shorter working period is:

$$P^* = H_0 / H_1$$

The three expressions can be read as the output identity, the saved-hours delta and the break-even productivity multiplier.

For example, if a worker previously worked 37 hours per week and a society or firm wanted to reduce the working week to 20 hours while maintaining the same output, the break-even productivity multiplier would be $37 / 20 = 1.85$. If productivity rises above that threshold, the model produces a surplus.

The model does not claim that every job can be compressed in this way. It does not account by itself for coordination costs, quality control, demand limits, fatigue, sectoral differences, regulation, care work, physical work or presence-bound services. It is a first-pass visibility equation. Its function is to identify whether a measurable time surplus exists before asking who captures it.

2.1 Saved-hours audit protocol

For the model to become more than a metaphor, a saved-hours claim should pass a basic audit test. The audit unit should be a task, workflow, team or service line, not an entire economy. The first step is to define the baseline output: what is being held constant, over what period, and at what quality threshold. The second step is to measure the old human hours required to produce that baseline. The third step is to measure the new human hours required after an AI or productivity intervention, including prompt time, supervision, verification, error correction, compliance work, training time and coordination costs.

The fourth step is the capture test. If hours fall while output remains constant, the audit should record where the difference goes: shorter required working time, higher wages, lower staffing, more output, shorter queues, lower prices, higher margins, quality improvement, more documentation, more monitoring or denser work. The fifth step is the metric-capture test: did the definition of acceptable output change after the productivity gain became available?

A saved-hours audit therefore has six minimum fields:

Audit field	Test question
Baseline output	What output is held constant?
Quality floor	What error, review, safety or service threshold must not fall?
Old human hours	How many human hours were required before the intervention?
New human hours	How many human hours are required after all AI-related supervision and correction work?

Capture channel
Output redefinition

What did the difference become?
Did the institution raise the output
standard after the gain appeared?

In a hypothetical public casework audit, a team previously spent 600 staff hours per month preparing standard case summaries for a stable caseload. If AI-assisted drafting reduces preparation time to 420 hours, but review, correction and appeal-risk checks add 60 hours, the net saved-hours estimate is 120 hours. That is the technical surplus. The political question is whether those 120 hours become shorter working time, faster service, more cases, more documentation, lower staffing, or merely a new expectation that each worker handles a larger monthly caseload.

3. Definitions

Saved hours are human labour time made technically unnecessary for maintaining a defined baseline of output after a demonstrated productivity increase.

A time dividend is the portion of saved hours converted into reduced required working time for humans.

Capture is the institutional allocation of saved hours to time, wages, output, profit, prices, state capacity, displacement, service quality or work intensity.

Metric capture occurs when saved hours disappear because the definition of acceptable output expands after productivity rises.

These distinctions matter because saved hours can exist without becoming free time. A firm may use the gain to produce more. A state may use it to process more cases. A platform may use it to raise targets. A worker may receive higher wages but not fewer hours. A consumer may receive lower prices. A public service may reduce waiting time. All of these are possible capture pathways.

The time dividend is narrower. It exists only when productivity gains are converted into less required human labour time.

4. Output Definition and Metric Capture

The saved-hours model depends on a baseline. But in many modern jobs, output is not fixed.

A factory can count units. A delivery platform can count deliveries. A call centre can count calls. But in law, care, teaching, software, administration, management, design, analysis and public service, output can be redefined.

If AI reduces the time required for a task, management may expand the definition of acceptable output. A legal review becomes legal review plus deeper

risk mapping. A lesson plan becomes lesson planning plus individualised feedback. A public case file becomes case file plus additional documentation. A software task becomes software plus more tests, edge cases, reporting and features.

In such cases, AI may save time at the task level while the institution captures that time by redefining output.

Measurement is therefore not neutral. Whoever defines output can capture saved time before it becomes negotiable.

5. Uneven Compressibility

Saved-hours analysis must distinguish between compressible, assistable and presence-bound work.

Compressible tasks are tasks where AI can reduce human time without fundamentally changing the task: routine drafting, translation, summarisation, extraction, ticket classification, report preparation, standard emails and first-pass code.

Assistable but bounded tasks are tasks where AI may reduce surrounding work while the core activity remains relational, judgement-based or human: medical consultations, teaching, therapy, social work, legal advice, public casework and management conversations.

Presence-bound tasks are tied to physical presence, safety, waiting, care, monitoring or maintenance: nursing presence, emergency response, construction, transport, inspection, childcare, eldercare, infrastructure operations and night shifts.

The saved-hours doctrine applies most cleanly to tasks and workflows, not whole jobs. A doctor may save time on notes, but that does not mean the patient encounter should become five times shorter. A social worker may save time on forms, but trust cannot be automated into existence.

6. Productivity Risk

The model assumes that productivity gains can be measured. In practice, AI productivity gains are uneven and uncertain.

Some tasks become much faster. Some barely improve. Some require new verification work. Some produce errors humans must correct. Some create compliance burdens. Some shift time from production to supervision. Some raise output expectations faster than they reduce labour.

Saved-hours policy therefore requires risk allocation. If a firm assumes a 30 percent productivity gain and cuts staff before the gain is stable, the downside may fall on remaining workers as stress, backlog and overtime. If a public authority assumes that AI will reduce administrative workload and

cuts budgets too early, the downside may fall on citizens as delays, errors or automated rejection.

The point is not to deny productivity gains. Experimental and field evidence already shows that generative AI can raise measured productivity in some writing and customer-support settings (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2023). The point is that task-level productivity is not yet a time dividend. A measured gain becomes a social fact only after the new supervision burden, error burden and output baseline are known.

Where productivity gains are uncertain, the policy question is not only how to share the upside. It is who absorbs the downside.

7. Capture Framework

The saved-hours framework distinguishes several capture models:

Capture model	What the productivity gain becomes
Worker time capture	Reduced required working time with stable income or stable output.
Firm output capture	Same hours, more production, faster delivery or higher throughput.
Capital capture	Higher margins, valuation, shareholder return or reduced labour cost.
Consumer capture	Lower prices, wider access or better service.
State capacity capture	Administrative, fiscal, military, welfare or surveillance capacity.
State-coordinated distribution	Industrial upgrading, social policy, public goals or national strategy.
State-mediated productivity bargain	Wages, skills, competitiveness and institutional mediation.
Measurement without distribution	Output per hour rises, but no institution converts the gain into time.
Social-dialogue capture	The gain is bargained into time, wages, staffing, flexibility or quality.
Metric capture	Output standards rise, hiding saved hours before distribution.
System intensity capture	Work becomes denser, more monitored and less slack.
Reverse time dividend	Productivity pressure is answered by more available working time rather than less.
Mixed capture	The gain is distributed across several channels.

The least visible model may be system intensity capture. Workers keep the job and the wage, but the pauses inside the job disappear.

8. Case Tests

8.1 China: productivity as state and industrial capacity

China is a useful stress test because its policy language around technology does not frame AI only as firm-level efficiency. The concept of new quality productive forces links innovation, high technology, high efficiency and high-quality development to national transformation.

For the saved-hours doctrine, the question is not whether China uses this language. It is what capture pathway this language implies. If AI or related technologies create new productive capacity, does the gain become shorter working time, higher wages, industrial upgrading, state capacity, platform control, national competitiveness or deeper labour pressure?

China's 2024 national time-use survey also shows why saved hours cannot be analysed only inside employment contracts. Paid labour, unpaid labour, training, ICT use, transport and discretionary activity form a broader time economy. A reduction in formal working time could become free time, but it could also become unpaid care, platform work, second jobs, training pressure, household labour or digital availability.

The case test is therefore institutional rather than causal. Evidence that a state frames AI and advanced productivity as national transformation is not evidence that AI has already saved a measurable number of hours for workers. It is evidence about the likely direction of capture if a surplus appears. The expected capture channel is industrial upgrading and state capacity first; a time dividend would require separate evidence that working-time obligations or workload baselines were reduced.

China therefore tests productivity as state and industrial capacity rather than automatic human time.

8.2 Singapore: productivity as wage and skills bargain

Singapore is a different stress test. It is a coordinated productivity state where government, employers and labour institutions often work through tripartite frameworks.

Singapore's National AI Strategy 2.0 frames AI as a force for public good, national capability, people-centred deployment and broad adoption. Its Progressive Wage Model links wage progression to skills upgrading and productivity improvement.

This is not a saved-hours dividend. It is a productivity-to-wage pathway.

Singapore shows that productivity gains can be institutionally mediated without becoming shorter working time. The gain may be channelled into wages, workforce upgrading, national competitiveness, business transformation, public-sector capability or social stability.

The case is important because it prevents the doctrine from treating all non-time outcomes as failure. A productivity gain that becomes wages, skills or public capability may be legitimate. The saved-hours question is narrower: did the institution explicitly decide that the gain should not become reduced hours, or did the possibility of reduced hours never enter the bargain?

8.3 United States: measurement without distribution

The United States is a measurement case. The Bureau of Labor Statistics defines labour productivity as output per hour worked (BLS Handbook of Methods). That statistical grammar already makes the saved-hours question legible.

If output per hour rises, the distribution question follows: does the gain become shorter working time, higher wages, lower prices, higher profit, more output, lower headcount or denser work?

BLS also notes that AI is difficult to isolate directly in productivity statistics and is captured indirectly through software, capital input, intellectual property products and productivity accounts (BLS, 2026). That supports a methodological caution: saved-hours claims should be audited against actual output, hours, quality and risk before they become political claims.

Task-level evidence makes the measurement gap sharper. Studies of generative AI in professional writing and customer support report productivity improvements in bounded settings (Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2023), but those studies do not by themselves answer the distribution question. A call-centre productivity gain can become fewer worker hours, more resolved tickets, tighter performance targets, lower service costs or higher margins. A writing-task gain can become earlier finish times, more documents, more review layers or higher expectations for polish.

The United States therefore shows that measurement can identify productivity change without deciding who receives the saved hour. It also shows why saved-hours audits would need to sit between productivity statistics and labour bargaining. Official productivity accounts can show that output per hour has changed; task-level experiments can show that particular workflows became faster; neither source, by itself, identifies whether the resulting time surplus was returned to workers, absorbed by output expectations or captured by capital.

8.4 European Union: social dialogue and algorithmic management

Europe is the institutional bargaining case. Working time, social dialogue, platform work, algorithmic management and AI governance already sit in the same policy field. The EU Working Time Directive makes working time, rest and maximum weekly hours legal governance objects (Directive 2003/88/EC). The AI Act places certain employment, worker-management and access-to-

work AI systems within the high-risk architecture, including systems connected to recruitment, selection, work-related decisions, task allocation, monitoring and evaluation (Regulation (EU) 2024/1689). Neither instrument creates a saved-hours dividend, but together they show that Europe already has legal categories for both time and AI-mediated workplace control.

European Commission survey material indicates that approximately one in three EU workers use AI tools at work, especially for text-related tasks (European Commission, 2025), while digital monitoring and automated scheduling are also part of workplace digitalisation. Eurofound's work on working-time reduction and the four-day week shows that reduced hours, flexibility and compressed schedules are active European debates. Autonomy's *GPT-4 (Day Week)* reports make the AI-working-time link more explicit by modelling scenarios where AI-enhanced productivity is captured as reductions in weekly working hours with pay protected (Garcia, Kikuchi and Stronge, 2023).

This makes Europe a contested case rather than a single model. Social dialogue is the bargaining pathway: it can make saved hours visible enough to negotiate over time, wages, staffing and workload. Algorithmic management is the absorption pathway: it can capture potential saved hours by monitoring, organising, directing and evaluating work through automated or semi-automated systems before bargaining is triggered.

The policy problem is that the two mechanisms operate at different speeds. Bargaining tends to require evidence, representation, procedure and time. Algorithmic management can change targets, scheduling, routing, evaluation and workload density inside the operational system itself. If saved hours appear first inside dashboards and workflow software, they may be absorbed before they become a collective bargaining object.

On the evidence used here, both mechanisms are active. The paper does not claim which dominates across Europe. It treats Europe as the policy field where the contest between bargaining and absorption is most explicit. A European saved-hours audit would therefore need to connect working-time law, AI-system transparency, workplace consultation and workload-density measurement.

8.5 Greece: the reverse case

Greece tests the doctrine in reverse. Where productivity is structurally weak, investment and technology adoption lag, and labour-supply pressure is high, one possible institutional response is to make more working time available rather than less.

Greek labour rules distinguish contractual working time from wider statutory and overtime limits. The standard contractual working week is 40 hours, while statutory and overtime structures allow additional work under defined conditions. In 2024, the activation of sixth-day work registration through ERGANI became a European working-time controversy.

The measure was contested by labour organisations and debated internationally because critics saw it as weakening the five-day norm, while the government framed it as a limited and registered measure for specific operational conditions. The important legal boundary is that Greece did not simply introduce a universal six-day week. It activated a framework in which sixth-day work can be registered and used under specified conditions.

OECD analysis of Greece points to low productivity, investment gaps and weak technology adoption as structural problems. The saved-hours implication is not that the sixth-day measure was caused by AI, or by productivity weakness alone. It is that time policy can move in the opposite direction when productivity gains are weak, uneven, not distributed or not institutionally converted into reduced required labour time.

When productivity does not free time, institutions may ask humans to supply more of it.

The reverse case therefore does not prove a mechanical law. It marks a warning boundary. A country, firm or public service facing capacity pressure can answer the productivity problem in two broad ways: make each hour produce more, or make more hours available. The saved-hours doctrine is concerned with the first path because it creates a distribution question. Greece illustrates the second path because it shows that institutional direction can move toward time extension rather than time release.

8.6 Denmark: measurement and public-sector comparator

Denmark is a comparator, not the main European case. Statistics Denmark's Working Time Accounts provide national measurement infrastructure around hours worked. Danish AI strategy material frames AI as a tool for public-sector improvement, productivity and responsible deployment with people and democratic values in view.

Denmark shows that strong measurement and public-sector digital capacity can support a saved-hours audit. It does not by itself guarantee a time dividend.

The Danish comparator matters because it separates capacity from allocation. A state can have high-quality labour statistics, strong digital administration and public-sector AI ambitions while still lacking a formal rule that converts productivity gains into reduced required working time. In a Danish setting, a saved-hours audit would be technically plausible because the measurement infrastructure is strong. The unresolved question would be institutional: would documented saved hours become shorter time, better service, staffing relief, more output or budget savings?

9. Policy Implications

First, governments and labour institutions should develop saved-hours audits. These would estimate how many hours AI or other productivity tools save in specific tasks, workflows or sectors at constant output and quality, and would mark uncertainty where gains are not yet demonstrated.

Second, productivity-sharing agreements should define how documented gains are allocated across time, wages, staffing, output, prices, profit and service quality.

Third, shorter-week pilots should be linked to actual output, quality, error rates, stress, retention and workload displacement, not only employee satisfaction.

Fourth, public-sector AI should be evaluated by burden reduction as well as throughput. A state can process more cases while making life heavier for citizens and public employees.

Fifth, anti-intensification safeguards should measure workload density, monitoring, cognitive pressure and autonomy. If every saved minute is immediately refilled, saved hours remain technical but never become human time.

Sixth, workers and their representatives need mechanisms to contest output redefinition. Without such mechanisms, metric capture can absorb saved hours before they become visible.

Seventh, educational and public-comprehension work should test whether the saved-hours question can be understood outside expert policy language. A concept that cannot be explained clearly to young people, workers or citizens is unlikely to become a usable democratic policy object. For this reason, a small school-based pilot on saved-time intuitions may be useful as a public-language test, though not as labour-market evidence.

The practical policy question can be stated as a capture matrix:

If saved hours are captured as...	Required policy question
More output	Who approved the new output baseline, and was workload density measured?
Higher profit or lower labour cost	Were workers compensated for the productivity gain?
Shorter working time	Was income, staffing and service quality protected?
Public-sector throughput	Did burden fall for citizens and workers, or merely move elsewhere?
Lower prices	Was the price gain traceable to productivity rather than degraded work?
Higher quality	Was the quality standard explicit before the gain was claimed?

More monitoring

Did autonomy, stress or discretion worsen?

Lower headcount

Was the productivity gain stable enough to absorb downside risk?

10. Counterarguments and Failure Modes

The first objection is that saved hours cannot be measured cleanly because output changes. This is partly correct. The doctrine does not assume that output is naturally fixed. It treats output definition as the central governance problem. If the output baseline cannot be defined, the institution should not claim a clean saved-hours gain.

The second objection is that AI creates new work. This is also correct. Verification, supervision, compliance, data governance, prompt design, exception handling and model-risk management can consume part or all of the apparent gain. That is why the audit counts net saved hours, not gross task acceleration.

The third objection is that productivity gains need not become free time. They may reasonably become better service, lower prices, higher wages, public capacity or investment. The doctrine does not rank time as the only legitimate use of productivity. It argues that time should be visible as one possible claim on the surplus.

The fourth objection is that labour shortages make reduced hours irresponsible in some sectors. In care, health, public administration, infrastructure and education, saved hours may be better used to reduce backlogs, increase presence or improve quality. That is a capture decision, not a reason to leave the saved hour unmeasured.

The fifth objection is that a shorter formal week can hide intensification. A four-day week that preserves pay but compresses five days of stress into four may be a schedule change rather than a time dividend. A real time dividend requires attention to workload density, autonomy, pace, monitoring and recovery time.

One further failure mode is historical and political rather than technical: autonomy anxiety can harden into anti-technological absolutism. One route treats large technical systems as totalising captures of human agency and answers with withdrawal, sabotage or destruction. The opposite route treats technological power as governable only if institutions can name its effects, protect human dignity, contest unjust use and distribute its gains. This paper rejects anti-technological absolutism and takes the second route. The saved-hours doctrine points toward institutional accounting, bargaining and distribution: measure the surplus, name the capture channel and make the allocation of time negotiable before unmeasured autonomy loss becomes a political vacuum.

These objections do not defeat the saved-hours doctrine. They define its measurement boundary. A weak saved-hours claim says "AI saves time." A stronger claim says where, how much, at what quality level, after what new work, and for whom.

11. Limitations

This paper proposes a conceptual and policy framework. It does not provide a full econometric model of AI productivity. It does not estimate sector-by-sector saved hours. It does not claim that all jobs can be compressed or that task compression equals whole-job compression.

The model is strongest for tasks and workflows where output can be defined and quality can be checked. It is weaker where output is relational, qualitative, presence-bound or politically contested.

The country cases are stress tests, not comprehensive national studies. They are selected to illustrate distinct institutional capture pathways: state capacity, wage bargaining, measurement without distribution, social dialogue, algorithmic management and reverse time pressure. They should not be read as a statistically representative sample of global AI productivity governance.

The paper also does not claim that shorter working time follows automatically from AI. That is the point. Saved hours become free time only through institutional choice.

This version also identifies two post-V0.92 development routes rather than treating them as completed evidence. First, *Magnifica Humanitas* provides a new social-doctrine context for thinking about AI, work, dignity and technological power. Second, a proposed Danish school pilot may test whether pupils can distinguish saved time from returned time, more output and intensified expectations. Neither route proves the labour-market claims in this paper. They strengthen the public framing and future research path.

12. Conclusion

The future of work should not be measured only in jobs lost or jobs saved. It should also be measured in hours saved, hours captured and hours returned.

AI may create productivity gains. But productivity is not distribution. A society can become richer without becoming freer if productivity is not converted into human time.

The saved-hours doctrine therefore asks a prior question: before a demonstrated productivity surplus disappears into profit, output, state capacity, metrics or work intensity, can it be made visible?

If AI governance needs a jobs policy, it also needs a time policy.

The deeper issue is not only whether AI improves productivity. It is whether a society can govern productivity without losing sight of the person whose time is being reorganised. The hour is not just an input. It is a piece of a human life.

References

- International Labour Organization. **C1 Hours of Work (Industry) Convention, 1919.** 28 November 1919. URL: www.ilo.org/resource/c1-hours-wor... Accessed 21 May 2026.
- International Labour Organization. **Convention No 1: A landmark for Workers' Rights.** 13 November 1919. URL: www.ilo.org/resource/article/conv... Accessed 21 May 2026.
- Mill, John Stuart. **Principles of Political Economy, Book IV, Chapter VI.** Ashley edition, 1909. URL: oll.libertyfund.org/titles/mill-p... Accessed 21 May 2026.
- Marx, Karl. **Grundrisse: Foundations of the Critique of Political Economy.** Translated by Martin Nicolaus, Penguin Books, 1973. Fragment on Machines / Notebook VII used as bounded historical framing.
- **Industrial Society and Its Future.** 1995. Boundary-case source only; used as a rejected example of anti-technological absolutism and autonomy-loss framing, not as authority, endorsement or policy support. URL: en.wikisource.org/wiki/Industrial... Accessed 24 May 2026.
- Xinhua. **Explainer: What do "new productive forces" mean?** 21 February 2024. URL: english.news.cn/20240221/3e0d1b79... Accessed 21 May 2026.
- The State Council of the People's Republic of China / Xinhua. **China unleashes new quality productive forces in push for reform, innovation.** Updated 25 June 2024. URL: english.www.gov.cn/news/202406/25... Accessed 21 May 2026.
- National Bureau of Statistics of China. **Communique on China's Third National Time Use Survey (No. 2).** 31 October 2024; page timestamp 1 November 2024. URL: www.stats.gov.cn/english/PressRel... Accessed 21 May 2026.
- Singapore Ministry of Manpower. **Progressive Wage Model.** URL: www.mom.gov.sg/employment-practic... Accessed 21 May 2026.
- Smart Nation Singapore. **National AI Strategy 2.0.** 2023. URL: file.go.gov.sg/nais2023.pdf. Accessed 21 May 2026.
- U.S. Bureau of Labor Statistics. **Productivity and Artificial Intelligence.** Last modified 4 May 2026. URL: www.bls.gov/productivity/articles... Accessed 21 May 2026.

- U.S. Bureau of Labor Statistics. **Handbook of Methods: Productivity Measures: Concepts.** URL: www.bls.gov/opub/hom/opt/concepts... Accessed 21 May 2026.
- Noy, Shakked, and Whitney Zhang. **Experimental evidence on the productivity effects of generative artificial intelligence.** *Science* 381, no. 6654 (2023): 187-192. DOI: doi.org/10.1126/science.adh2586.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. **Generative AI at Work.** NBER Working Paper No. 31161, 2023. DOI: doi.org/10.3386/w31161.
- European Commission. **Survey results show that one in three EU workers use AI tools at work.** 21 October 2025. URL: employment-social-affairs.ec.europa.eu/en/employment-social-affairs/2025-10-21-survey-results-show-that-one-in-three-eu-workers-use-ai-tools-at-work Accessed 21 May 2026.
- European Parliament and Council. **Directive 2003/88/EC concerning certain aspects of the organisation of working time.** 4 November 2003. URL: [eur-lex.europa.eu/legal-content/E...](https://eur-lex.europa.eu/legal-content/EU/summary/dir/2003/88/eng.html) Accessed 22 May 2026.
- European Parliament and Council. **Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).** 13 June 2024. URL: [eur-lex.europa.eu/legal-content/E...](https://eur-lex.europa.eu/legal-content/EU/summary/reg/2024/1689/eng.html) Accessed 22 May 2026.
- Eurofound. **Four days a week? Europe debates shorter working times.** Published 6 October 2025. URL: [www.eurofound.europa.eu/en/public...](https://www.eurofound.europa.eu/en/publications/2025-10-06-four-days-a-week-europe-debates-shorter-working-times) Accessed 21 May 2026.
- Eurofound. **Working time reduction with a focus on the four-day week: Literature review.** 2024. URL: [www.eurofound.europa.eu/en/public...](https://www.eurofound.europa.eu/en/publications/2024-04-04-working-time-reduction-with-a-focus-on-the-four-day-week-literature-review) Accessed 21 May 2026.
- Garcia, Luiz, Lukas Kikuchi, and Will Stronge. **GPT-4 (Day Week): Great Britain Edition: How the latest AI technology can lead to a nation of shorter hours, without reductions in livelihoods or productivity.** *Autonomy*, 20 November 2023. URL: [autonomy.work/portfolio/gpt-4-day...](https://autonomy.work/portfolio/gpt-4-day-week-great-britain-edition) PDF: [autonomy.work/wp-content/uploads/...](https://autonomy.work/wp-content/uploads/2023/11/gpt-4-day-week-great-britain-edition.pdf) Accessed 26 May 2026.
- Garcia, Luiz, Lukas Kikuchi, and Will Stronge. **GPT-4 (Day Week): US Edition: How the latest AI technology can lead to a nation of shorter hours, without reductions in livelihoods or productivity.** *Autonomy*, 20 November 2023. URL: [autonomy.work/portfolio/gpt-4-day...](https://autonomy.work/portfolio/gpt-4-day-week-us-edition) PDF: [autonomy.work/wp-content/uploads/...](https://autonomy.work/wp-content/uploads/2023/11/gpt-4-day-week-us-edition.pdf) Accessed 26 May 2026.
- Eurofound. **Platform work: Algorithmic management.** URL: [www.eurofound.europa.eu/en/platfo...](https://www.eurofound.europa.eu/en/platform-work/algorithmic-management) Accessed 21 May 2026.

- International Labour Organization. **Working Time and Work-Life Balance Around the World.** 6 January 2023. URL: www.ilo.org/publications/working-... Accessed 22 May 2026.
- Hellenic Labour Inspectorate. **Limits on working time.** URL: www.hli.gov.gr/en/ergasiakes-sche... Accessed 21 May 2026.
- Greek Ministry of Labour and Social Affairs. **Digital work card and sixth-day work registration note.** 21 June 2024. URL: ypergasias.gov.gr/ti-ischyei-me-t... Accessed 21 May 2026.
- OECD. **OECD Economic Surveys: Greece 2024: Boosting firm growth and innovation.** 2024. URL: www.oecd.org/en/publications/oecd... Accessed 21 May 2026.
- Statistics Denmark. **Working Time Accounts.** URL: www.dst.dk/en/Statistik/emner/arb... Accessed 21 May 2026.
- Agency for Digital Government. **Danish strategies for artificial intelligence.** URL: en.digst.dk/digital-governance/ne... Accessed 21 May 2026.
- Pope Leo XIV. **Magnifica Humanitas: On Safeguarding the Human Person in the Time of Artificial Intelligence.** 15 May 2026. URL: www.vatican.va/content/leo-xiv/en... Accessed 26 May 2026. Used as post-V0.92 social-doctrine context, not as evidence for the saved-hours model.

Appendix Candidate: School Pilot On Saved-Time Intuitions

This appendix is a placeholder for a possible small Danish school pilot. The pilot would not test the labour-market claims in this paper. Its value would be different: to test whether the saved-hours question is understandable as a public question.

The core scenario is deliberately simple:

A school or workplace gets a new tool. A task that used to take 10 hours now takes 6 hours. The result is just as good. Four hours have been saved. What should happen to the four hours?

If pupils can distinguish saved time from returned time, more output, higher pressure and shared decision-making, the concept may have public-policy clarity beyond expert labour debates.

The pilot should remain anonymous, minimal and low-risk. It should collect no names, no sensitive personal data and no school-level ranking. Results, if collected, should be described as exploratory public-comprehension material, not as representative opinion polling or evidence of economic outcomes.