

We Saved Eight Hours. Did We?

Why measuring saved human time turned out to need a protocol.

DRAFT PREVIEW DENNIS HEDEGREEN OPEN DRAFT-LOCAL

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/WE-SAVED-EIGHT-HOURS-DID-WE/](https://hedegreenresearch.com/articles/we-saved-eight-hours-did-we/)

Why measuring saved human time turned out to need a protocol

Suppose something used to require ten hours of human work and now requires two. The intuitive conclusion is straightforward: eight hours have been saved. That simple subtraction was roughly where this research started.

A few months ago I wrote *The Saved Hours Doctrine* around a question that seemed more interesting to me than the usual productivity discussion. If technology allows us to maintain the same output with less human labour, what actually happens to the human time that is no longer technically required? It might become shorter working time, additional output, higher profit, lower prices, fewer workers, greater public capacity, or entirely new tasks. A productivity gain by itself does not tell us which of those outcomes follows.

That was why I initially treated Saved Hours as a downstream problem. I was interested in who captures the released time and what institutions determine where it goes. I assumed that the upstream measurement problem would eventually be attacked by people who already work inside labour economics, industrial engineering, work measurement, productivity research and trade unions. I expected specialists to tell me where the idea was wrong, which parts were already solved, and what a serious measurement protocol would have to look like.

But the more I tried to use Saved Hours, the harder it became to avoid a more basic problem. Before asking who captured eight saved hours, I needed some defensible way to establish that eight hours had actually been saved.

That turned out to be much harder than the subtraction.

Ten hours divided by two is still easy

If one production state requires ten human hours and another requires two for the same output, the simplest comparison is

$$LF = H_0 / H_1 = 5$$

I call that ratio **Labourfold**. The first state requires five times as much human person-time as the second. The arithmetic itself is ordinary labour-productivity arithmetic expressed from the human-time side. The problem is deciding whether the two numbers are actually comparable.

Suppose the ten-hour process produces 100 units while the two-hour process produces only 10. Ten divided by two is still five, but it no longer tells us that the same output requires one fifth as much human labour. The calculation is correct; the claim is not.

The arithmetic is correct. The claim is not.

The same problem appears with quality. If two systems both produce 100 units, but one produces 99 acceptable outputs and the other produces 70, the apparent saving may simply come from doing a worse job. The relevant comparison is therefore not just “how much was produced?” but whether both states still provide the same material service.

AI makes this especially easy to see. A model may generate a draft in seconds, but those seconds are not automatically the new human-time requirement. If a person still has to read the result, verify facts, correct mistakes, resolve exceptions or take responsibility for the final output, that required human work has not disappeared simply because generation became fast.

The same applies when labour moves rather than vanishes. A supermarket can reduce cashier hours by moving part of the checkout process to customers. That may be a real organisational saving for the supermarket, but it is not automatically an equivalent reduction in the total human time required to provide the service. The same issue appears when work moves to contractors, vendors, patients, citizens or remote operators. A change in who performs a necessary human function is not the same thing as removing the function.

Standby creates another version of the problem. A worker may go from actively operating a machine for four hours to intervening for only one, but if that person must still remain available for the entire eight-hour shift, the system has made them less busy without necessarily releasing their time.

Less busy is not the same as free.

That is where Labourfold stopped being a division problem and became a measurement protocol.

Sometimes the correct answer is no number

You can almost always divide two numbers. That does not mean you should.

If the output has changed too much, if the service quality is not comparable, if human work has merely moved outside the visible organisation, if the production states are materially different, or if the evidence does not actually identify required human time, then a neat scalar can create more confidence than the evidence deserves.

For that reason, Labourfold is allowed to refuse the calculation. The protocol can return a canonical Labourfold when the comparison is strong enough, a proxy ratio when the evidence is weaker but still useful, a zero-human-time endpoint when the declared target genuinely reaches zero required human time, or **no valid scalar** when the comparison fails materially.

That refusal matters to me because the temptation in productivity debates is usually to force a number out of whatever data are available. Labourfold instead treats refusal as a legitimate result. A measurement protocol should sometimes tell us that the evidence does not earn the number we want.

I did not plan to build this

I did not start the Saved Hours work because I wanted to design a new measurement protocol. I expected other people to attack the problem. I still do.

I am not an industrial engineer, labour economist, metrologist or union research department. There are people who understand individual parts of this problem far better than I do. My original hope was that Saved Hours would provoke exactly those people into saying: this boundary is wrong, this kind of hour cannot be counted that way, this part already has a better method, or this comparison cannot be made at all.

But science does not become correct because the right person says it, and it does not become untestable because the wrong person asks the question. So I kept following the failure points.

The result is not a claim that I invented labour productivity arithmetic. I did not. Labourfold is a protocol built around a familiar ratio. It asks what the output is, what quality still counts as the same service, which human work belongs inside the boundary, which production states are actually being compared, what time horizon matters, and what kind of evidence supports the human-time figures.

The point is not to turn every technology into one headline number. The point is to make it harder to claim that human labour disappeared when the comparison actually shows something weaker, narrower or different.

The research started downstream

The odd part is that the research series developed in almost the opposite order from the logic it now has.

First came **Saved Hours**, which asked what happens to human time when technology makes some labour technically unnecessary. Who captures it? Does it become shorter working time, greater output, lower prices, higher profit, fewer workers, new tasks or something else?

Then came *Post-Work Requires Power*, the Japan field paper. That paper starts one step later again. If some released time actually reaches human life, what material and energy access is needed to live inside it? A saved hour does not happen in empty space. People use time in homes, transport, care, learning, public services, workshops, networks, compute and all the other systems that make ordinary life possible.

Labourfold came after both of them and moved the research upstream. Before asking who receives the saved hours, and before asking what infrastructure those hours require, I still needed a disciplined answer to the simplest question: were the hours really saved?

So the historical order was:

Saved Hours → Post-Work Requires Power → Labourfold

But the logical order now looks cleaner:

Labourfold asks whether required human time actually fell for a comparable output and service.

Saved Hours asks what happens to the technical time release if it did.

Post-Work Requires Power asks what material access is needed if some of that released time is actually returned to people.

I did not plan that architecture. One question just kept opening the next one.

Attack it

I do not want Labourfold treated as something people are supposed to accept because I published a paper and attached equations to it. The useful outcome would be the opposite.

Attack it. Debunk it. Improve it.

If an existing discipline has a better rule for one part, replace mine. If a boundary is wrong, show why. If the evidence classes fail in a real domain, break them. If required human time cannot be identified in the way the protocol claims, demonstrate the failure. If unions, engineers, economists, firms, regulators or researchers can build a better protocol from the same question, that is progress.

The point is not to protect Labourfold. The point is to make claims about saved human time precise enough that disagreement can finally attack something concrete.

I am not finished asking questions

We still argue about working hours, retirement ages, labour shortages, productivity, automation and whether AI will remove jobs or create them. Those arguments matter, but underneath them sits a more basic question that we still measure surprisingly badly: **how much human labour is actually required to produce the things and services we want?**

If that requirement changes substantially, then a second layer of questions follows. Who gets the released time? What happens to wages, ownership and bargaining power? Which institutions were built around assumptions about work that may no longer hold? What infrastructure does a life with more non-work time require? What happens if some forms of necessary labour disappear much faster than others? And what does retirement mean if the quantity and structure of required human labour look very different from today?

I do not know the answers, and Labourfold does not answer them. It is only an attempt to make one upstream measurement problem less vague so the next questions can be asked more honestly.

That is enough for now.

I am not finished asking questions.

Read the research

The Labour That Output No Longer Requires *A Protocol for Measuring Required Human Time and the Saved Hours Question*

Permanent archive: <https://zenodo.org/records/23002321> (<https://zenodo.org/records/23002321>).

Coding Manual v1.6 is included in the same record.

The Saved Hours Doctrine and **Post-Work Requires Power** are part of the same developing research series.

RELATION MEMORY

EXTENDS [Where Did the Saved Hours Go?](#) · Starts from the Saved Hours question and moves upstream to measurement/admissibility.

NEARBY [The Saved Hour Also Needs Power](#) · Places Labourfold before the Saved Hours allocation question and the Post-Work material-access question.

NEARBY [What If the Slop Is Mine?](#) · Both articles treat apparent AI time savings as verification and responsibility problems.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/WE-SAVED-EIGHT-HOURS-DID-WE/](https://hedegreenresearch.com/articles/we-saved-eight-hours-did-we/)

PDF VIEW GENERATED: 2026-09-29T09:18:58.809Z