


# The end of history and the last job

## AI and the future of work

Paul Dixon | 30 Sep 2026



Recently I asked Google’s NotebookLM to transcribe and summarise a 35-minute political debate on YouTube. Nine seconds later, it was done. I had not acquired a superhuman ability. I had acquired access to one. Such moments are becoming familiar, but what they mean for work is only beginning to emerge. When a machine can perform in seconds a task that once occupied a person for an hour, the question soon becomes who still needs to be paid to do it.

Producing a summary is only one part of a job. Lawyers, journalists and analysts do more than process information, and a quick answer still needs checking. Yet much of their day consists of tasks that AI can now perform faster and cheaper. If enough of those tasks disappear, the jobs built around them may follow. The risk is that the gains will accrue chiefly to those who own the systems, while the costs fall on those whose labour they replace. The economy could become richer while fewer people can earn a decent living. Democratic institutions should shape how these gains are shared before the benefits and costs become entrenched.

### The productivity shock

NotebookLM did more than summarise the debate. It connected the material to other sources I had supplied and offered to turn the collection into a mind map, slides, flashcards or a narrated video. Work that once required several tools, and perhaps several people, had become a sequence of clicks. AI can now help with a whole chain of office tasks: gathering information, comparing sources, drawing conclusions and presenting the result. As the tools improve and connect more easily, employers will have reason to reorganise whole workflows around them.

Critics sometimes describe large language models as “stochastic parrots”: systems that predict words without understanding concepts. Whether or not that description is fair, the practical question is whether LLMs can produce useful, reliable work at lower cost. A system need not think like a person to reduce demand for a person’s labour.

Consider a simple thought experiment. A business employs 100 office workers and introduces an AI system. Set its initial capacity at 1.0 on an index and suppose that it doubles each year. After seven years, the index reaches 128. This is not a forecast; it simply makes the logic visible.

Now suppose the business maintains its output while progressively replacing human labour with AI. Routine tasks go first; the remaining tasks require ever greater advances to automate. Figure 1 illustrates this pattern. In the example, roughly two workers remain after seven years, and employment gradually approaches one supervisor as AI capacity grows. The shape and pace of the curve are assumptions, not predictions.




Figure 1 An illustrative transition from human labour to AI

In this illustration, roughly 98 of the original 100 jobs have disappeared by year seven. Some savings may reach customers through lower prices, and some may finance expansion. But if owners keep most of the gains, a larger share of the business's income will flow to them rather than to employees. Displaced workers will also compete for the jobs that remain, putting pressure on wages.

Real businesses are more complicated. Tasks do not fit together neatly, errors are costly, and customers may demand more as prices fall. These complications will determine how much employment survives. They do not remove the incentive to substitute cheap machine capacity for expensive human time. Expertise may become abundant while control over the systems supplying it remains concentrated. This is what I mean by a transfer of "knowledge power". It could alter the balance between labour and capital.

## Jobs that last

The transition will be uneven. Many tasks performed on a screen are easier to automate than physical work in unpredictable surroundings. A plumber repairing a leak in an old house faces a different problem from an analyst summarising a report. Scaffolding, hairdressing and mountain rescue still demand dexterity and judgement in conditions that are hard to standardise. Other jobs endure because human contact is part of what people value. A patient may want a doctor to explain a diagnosis, or a family may want a person to care for an elderly relative. In teaching, therapy and hospitality, the relationship can matter as much as the service. Some work also derives its appeal from a human having done it. We watch an athlete for the achievement and listen to a musician partly for the person behind the performance. AI may compete in these markets without extinguishing demand for human skill.

None of this guarantees enough employment for everyone. People displaced from offices will seek work elsewhere. More applicants for jobs in care, hospitality or the arts could depress earnings even where machines cannot replace the workers. Management is not automatically safe either. If a company needs fewer employees, it may also need fewer people to recruit, supervise and coordinate them. Some people will oversee AI systems, but there is no reason to assume that these roles will match the number lost. Manual work has a reprieve only while robotics lags behind software. How long that lasts is uncertain.

AI may itself help engineers improve robots, and a machine need not perform every human task to threaten a particular occupation. A robot that erects scaffolding need not also perform surgery. Over time, employment could become concentrated in activities for which customers specifically want a human. The size of that market will depend on what people value and what they can afford.

## **Who benefits**

The economic consequences turn on ownership. A society in which most people share the returns from AI would experience automation very differently from one in which those returns flow to a small minority. Shares in publicly traded companies offer one route to participation, including through pensions. But ownership is unequal, and many people have little wealth to invest. Replacing wages with returns on capital therefore risks widening the gap between those who own productive assets and those who must sell their time.

We do not need a precise ten-year forecast to see the problem. If AI raises profits while reducing the wage bill, owners can prosper as workers struggle. Taxation, competition, welfare payments and lower prices may spread the benefits. Whether they do so sufficiently is a political question.

There is also a problem of demand. One business's wage bill helps sustain demand for the products of others. Unless lost earnings are replaced, falling employment could weaken spending across the economy. Governments could face a similar squeeze as receipts from wages and consumption fall just when more people need support. An economy may become capable of producing far more while leaving many people with less purchasing power. Producing more does not ensure that everyone can afford a share.

## **What work gives us**

A job provides more than income. It gives the week a rhythm, brings people into contact with others and offers a place in society. Even an unsatisfying job can provide a sense of being needed. Losing it can unsettle an entire life.

The COVID-19 pandemic showed how quickly familiar routines could dissolve. For many office workers, commuting and daily contact with colleagues vanished almost overnight. Mass displacement by AI would be a different shock: people could lose both the routines of work and the prospect of returning to them.

More leisure could be a considerable gain. But leisure created by prosperity feels different from time emptied by unemployment. A society built around paid work will need new ways to distribute income and sustain purpose if paid work becomes scarce. Technology may change faster than these institutions.

## **Reasons for doubt**

The strongest objection is that technological upheaval has often created new work as it destroyed old occupations. Cheaper production can expand demand, while new industries can generate jobs that were previously hard to imagine. AI could follow that pattern. But new jobs may not appear in the right places, require the skills displaced workers possess or pay comparable wages. If they consist largely of processing information on a computer, they may themselves be candidates for automation. What matters is the pace and scale of job creation relative to job loss.

A second objection is that progress may slow. Models could encounter limits in data, reliability or reasoning. Synthetic training data may help, but does not prove that improvement can continue indefinitely. Nor do impressive demonstrations show that AI can handle every messy, consequential task in a real workplace. These uncertainties make confident timetables unwise. Yet today's limitations

offer no lasting guarantee of safety. AI could displace many workers without surpassing every human ability.

Energy, water and computing capacity could also constrain adoption. Greater efficiency may reduce the resources needed for each task, while wider use increases total demand. Environmental costs belong in the political calculation alongside gains in productivity.

Perhaps machines will do the work and humans will flourish in other ways. This is the most attractive possibility. Yet it depends on people having both a share of the resulting wealth and worthwhile ways to spend their time. An abundance of machine labour does not, by itself, provide either.

## **A political choice**

The bleak scenario is not inevitable. AI systems are developed and deployed within rules that societies make. Property rights, taxation and competition help determine who owns productive assets and who benefits from them. Those rules can change. Leaving the transition to the market would still be a political choice: existing patterns of ownership would shape the distribution of a potentially extraordinary new source of wealth and power. By the time the social costs became unmistakable, the beneficiaries might have even greater influence over the rules.

The case for regulation should therefore begin with a clear purpose: to ensure that rising productivity improves life for the wider population. Governments will need to consider how the gains are shared, how workers are supported through displacement and where human responsibility should remain essential. The speed of deployment matters because people and institutions need time to adapt. Regulation need not eliminate rewards for innovation. Those who build useful systems should be able to prosper. The challenge is to reward success without allowing concentrated economic power to undermine livelihoods and fairness. International coordination will be hard. Governments will fear falling behind in an AI arms race, while companies will seek favourable jurisdictions. Early leaders may also attract more investment, reinforcing their advantage. Cooperation will require governments to overcome these pressures.

My nine-second summary of a 35-minute debate was a small convenience. Multiplied across millions of workplaces, such conveniences could change who earns a living and who holds power. Democratic governments should act before that power becomes entrenched, so that the gains from AI improve life for the many rather than the few.