

# The Great AI Transition

## Or: the end of history and the last job.

26 February 2026 – A short Essay on AI and the future of work.

Paul Dixon

Yesterday I sped-watched, transcribed and summarised a 35-minute YouTube political debate in 9 seconds. Howso? I did not develop a superhuman power – I simply used a free AI tool from Google, Notebook LM. If you don't believe this, watch the video ([link](#)). Many people are greatly underestimating the impact AI will have on labour and jobs. The effect of AI tool efficiency means that some knowledge tasks that took, until recently, 8 hours to do, can now be done in under 3 minutes – a 200x productivity improvement. Some have said AI LLMs only predict a sequence of words but does not understand – they are “stochastic parrots” (Bender et al., 2021). But it is now impossible to distinguish reliably between human and AI outputs (see Appendix for examples). I argue that this extraordinary increase in AI productivity will quickly lead to job destruction, dislocation, economic inequality and social unrest. To prevent this from happening, we must use our democratic institutions to regulate the development of AI to serve what humans want for their societies, and not leave it to the market to decide.

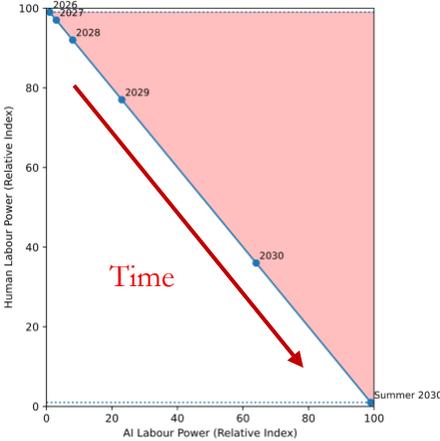
### A quantum leap in productivity gains

The transcription and summarisation capability I described above is making the jobs of journalists, lawyers, investment analysts, bankers, healthcare administrators, educational publishers, clerks, public sector office workers, call centre handlers, real estate administrators, customer service agents etc., mostly obsolete. Even allowing for frictional and management costs, this suggests the amount of labour needed for any task performed using a computer terminal or a smartphone has already been reduced by a factor of 100.

Not only did Notebook LM transcribe and summarise the 35-minute video in 9 seconds to a graduate-level standard, but it also linked it to every other source I had saved in this particular ‘context window’. A ‘one-click’ operation then created a mind-map, slide-deck, infographic, set of flashcards. It can create and narrate a video containing the merged content of *all* of the sources. AI never forgets. It could link this aggregation to every other aggregation in the world with which it shares content commonality. It can do this on every topic conceivable; from astrophysics to cancer drug development to small nuclear reactor manufacture, to warfare strategy. As this huge transfer of ‘knowledge power’ as I will call it occurs, the balance of power from humans to AI controllers shifts by an extraordinary amount. It may be more than you think; so let me explain.

Imagine you are observing a hypothetical small country at the start of the AI boom. For every 100 working humans there is 1 AI ‘engine’ in use, and it has a productivity equal to, let’s say, one office worker. Let’s call that power ‘1 Turing’. On a conservative measure, AI power is doubling every year. In 8 years, one AI engine will have the labour power of 128 office workers (consider the sequence 1, 2, 4, 8, 16, 32, 64, 128). Median estimates are faster, suggesting we will get there in 5 years. If an AI engine can soon do the work of 128 office workers, or let’s say, to be conservative, 100 office workers, we will have gone from a ratio of 100 workers to 1 AI Turing, to 1 worker for every 100 AI Turings. How so? Because the AI will make many of the 99 humans unemployed. The AI will do all the work previously done – without complaining, getting sick or needing to be hired and fired. They won’t go on paternity or maternity leave. The wages that went to labour now go wholly to the AI owners. One super-worker can ‘manage’ every 100-Turing-powered-AI – to the extent it needs to be managed. The remaining, singular employed human per hundred might be the smartest, most hardworking, or indeed the most political. However, with one hundred office workers competing for every job, the price of labour, by most conventional economic theories, will plummet. So not only does demand for office labour collapse, but so will wages. This scenario is described by Citrini & Shah (2026) in their speculative ‘report from 2028’:

“It should have been clear all along that a single GPU cluster in North Dakota generating the output previously attributed to 10,000 white-collar workers in midtown Manhattan is more economic pandemic than economic panacea. The velocity of money flatlined. The human-centric consumer economy, 70% of GDP at the time, withered. We probably could have figured this out sooner if we just asked how much money machines spend on discretionary goods. (Hint: it’s zero.) AI capabilities improved, companies needed fewer workers, white collar layoffs increased, displaced workers spent less, margin pressure pushed firms to invest more in AI, AI capabilities improved...It was a negative feedback loop with no natural brake. The *human intelligence displacement spiral* (Citrini & Shah, 2026).



**Figure 1: Office-worker labour to AI power, hypothetical projection over five years**

Description: By 2027, 3 workers in 100 have been displaced by AI with a power of 3 Turings. By 2029, 23 workers have been displaced. By 2030 64 have been displaced, and by Summer 2030 there is just one worker, 99 are unemployed, and there are 100 Turings of AI power doing all the office work.

## **Not everyone is unemployed....yet.**

Of course, I am only talking here about office workers. Some will initially be immune from replacement. First, **manual** workers, both skilled and unskilled. These include chefs, road-sweepers, plumbers, electricians, builders, hairdressers, hospital workers, physiotherapists, hotel cleaners, refuse collectors, farmers, soldiers, sports coaches. Any job that requires highly variable dexterity in challenging conditions is most protected – think of scaffolding or mountain rescue workers. Second, **‘must be human’ workers**, where interacting with a human is part of the service. This includes patient-facing doctors, psychiatrists, some teachers, carers, some police, personal trainers, waiters, some salespersons, commercial pilots, airline stewards and sex workers. Third, **creative arts workers**, where creation by a human or being a human is the role’s primary value. This includes actors, musicians, artists, a small number of writers – and football players. Furthermore, a handful of workers will be required to administer the AI systems. These will be employed by the owners of AI systems.

You will notice I have not mentioned ‘management’ or ‘leadership’ roles. These may comprise maybe up to 10% of the workforce today, but with AI doing up to 99% of office work in five or so years, and AI Robots doing much of the manual work in ten years, it is not clear how much human management will be needed. The human owners of AI will need some humans to ‘manage’ the AI systems for them, but conceivably human management and leadership will also be cut from 10% to maybe 2% of the workforce over ten years.

Manual workers will be protected for as long as robotics fails to keep pace with knowledge AI. However, given that ‘knowledge AI’ can be used to build robots, this might only take a few years to catch up. Drones are already replacing soldiers in warfare, for example. The remaining categories are, by definition, the hardest to replace – a new category of must-be-human, or ‘artisan’ workers describes these professions. Of course, development of a human-like robot that was sufficiently functionally equivalent might threaten these categories too. As office worker jobs disappear, some of the displaced will retrain as manual, must-be-human or creative workers, so pushing labour values down. The same process will repeat as manual workers are forced to compete for must-be-human or creative roles. Eventually, humans will only be employed in ‘must-be-human’ and creative work, limited by the demand for their services.

## **The economic effects of the transition**

Ownership of leading AI capabilities is remarkably concentrated in a few hands. OpenAI (Chat-GPT), Anthropic (Claude) and XAI (Grok) are essentially private companies. Google (Gemini) and Microsoft (Co-Pilot) are public, and so their ownership roughly tracks the public ownership of equity, although that too is quite unevenly distributed in populations. If we make a very general, idealised assumption that half of all AI ownership is held by the top 1% of the wealthy, and the other half by the next 20%, and we first imagine, unrealistically, that only computer-based white-collar work exists, then the wealth transfer from the labour displacement scenario above would see the top 1% gain in wealth 40 times, the top 20% gain in wealth by 2.5 times, and the remaining 80% decrease in wealth by about 90% (a

factor of 10). Now we adjust for the fact maybe only a third of labour is ‘office’ work, a third will remain manual, must-be-human or creative, and a third will transition over ten years, these numbers halve. Citrini & Shah (2026) put the proportion of white-collar workers in the U.S at 50%, so my estimate is quite conservative. But my example still results in the richest 1% gaining about 20 times; the next 20% (unevenly) gaining by 1.25 times, and the bottom 80% all declining in wealth by about 80%, by a factor of 5. This represents a huge increase in inequality, and even if it is tempered by taxation, social welfare, friction, cross-border complexities and inefficiencies, it will surely create significant political and social unrest. Further, consider if there are no wages, there will be vastly reduced consumer spending, and also much lower tax receipts from labour or consumption. This example is approximate and highly stylised, but you can probably agree it is not so unrealistic.

## The social effects of the transition

In addition to the effects of the economic inequalities just discussed, the social effects of disruption to work will be significant. Much of social structure is built around work practices, from the rhythm of commuting, working five days a week, to social bonds formed through working with other humans. All of this will rapidly be upended. The COVID-19 pandemic was a preview of this: knowledge workers stayed at home, but manual and must-be-human work continued. However the AI transition will be like a COVID shock with no end date. The mental health and long-term social effects of the COVID-19 pandemic are now well accepted, so we should not doubt the negative impact this would have. This is not to mention some of the negative effects of AI-driven tech itself, such as the political and mental health effects of social media, but that is for another day.

## Objections to the above scenario

1. **AI-LLM power will plateau.** *Response:* This objection has been raised many times, and every time AI exceeded the previously declared ‘ceiling’.
2. **AI-LLMs will have run out of words to consume.** *Response:* LLMs can generate ‘synthetic data’ to train on. They can recombine all known words to uncover new meanings and insights.
3. **We don’t have enough power, water, etc to run the AIs.** *Response:* Power and water consumption is falling per unit of compute power as technological advances are made. Of course, there is an environmental cost to AI that we may not want to bear – that is a political decision.
4. **AI-LLMs will never be able to master complex strategic tasks, exercise human judgement, creativity.** *Response:* AIs are already better at nearly all complex tasks than nearly all humans. Human judgement comes from experience, but AIs are rapidly gaining that too. Creativity is being seen in many new AI models such as those creating art or music.
5. **AI Robots will never be able to do what a human can do.** *Response:* This is a solvable engineering problem. It is not necessary that the same robot can both erect scaffolding and perform heart surgery. Robot specialisation is permitted.

6. **AI companies will be regulated and heavily taxed.** *Response:* Given global competitive (‘AI-race’) forces, there are strong political market-driven incentives against this and no evidence that this will happen anytime soon.
7. **Labour law and unionisation will prevent large scale redundancies.** *Response:* they may slow, but ultimately not prevent these outcomes.
8. **New jobs will be created.** *Response:* True, some new roles are emerging, mostly in AI prompt management, and also in technology delivery and security. However these are mostly ‘computer’ jobs that too will mostly be replaced in time.
9. **AI will do all the work, and humans can flourish in other ways.** *Response:* In 1930 Keynes predicted that by 2030, humans would only need to work 15 hours per week. He might be spot-on. Humans have developed leisure activities to fill the time freed up from labour, and the idea of vacation is a modern invention. But human nature and society is inherently not configured to live in a world of eternal leisure. Think of the “bored, wicked rich” (Hersh, 1998).

## What to do: why AI regulation would be a good thing

I have written above as if all were inevitable, but we are still, for now, in charge of how these systems are developed and deployed. It may be that we do not want it to end up this way. However, what makes the scenario a likely outcome, tempered maybe by violent political revolutions, is that the institutions of finance, law and politics are currently configured to permit the concentration of AI power in the hands of very few. The law protects the property rights of shareholders, and the machinery of finance and politics protects this. Governments, on the whole, have either chosen not to, or have been shut out of, or paid off from regulating the AI Technology sector. Now, for the past eighty years, this system of property rights, law and governance has worked out pretty well for the majority, but that was in a time of *relative* equality. We are presently returning to levels of financial and political inequality last seen in the ‘Gilded Age’ of the U.S., or Victorian Britain. Except this time, the differences are projected to be much greater and spread globally by technology’s reach. In democratic countries, we have a political choice, to regulate and direct how AI is deployed for the benefit of the widest range of humans. Regulation can still permit the brightest and the most hard-working to realise their dreams. There will still be tech and AI-powered billionaires. But we may be able to save the rest of the labour force from suffering the damaging, irreversible consequences of unregulated AI.

## References

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021).

On the dangers of stochastic parrots: Can language models be too big?  *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.  
<https://doi.org/10.1145/3442188.3445922>

Citrini, N., & Shah, A. (2026). *Report from 2028: The human intelligence displacement spiral* [Speculative scenario].

Hersh, K. (1998). *Like you* [Song]. On *Strange angels*. 4AD.

Keynes, J. M. (1930). Economic possibilities for our grandchildren. In *Essays in persuasion* (1963 ed.). W. W. Norton.

## **Appendix - Which was AI generated?**

### **Summary A**

In this interview, Microsoft AI CEO Mustafa Suleyman defends the unprecedented capital investment in artificial intelligence by highlighting the linear relationship between computing power and advanced capabilities. He shifts the industry focus from the abstract concept of General Intelligence toward "humanist superintelligence," envisioning highly autonomous, creative systems that remain strictly subordinate to human well-being and control. Suleyman outlines a future of AI self-sufficiency at Microsoft, where "digital minds" will soon handle complex professional-grade tasks and revolutionize fields like medicine through high-accuracy diagnostics. Ultimately, the discussion serves to balance technological acceleration with a warning against granting models "rights," emphasizing that these tools must be designed to enhance rather than exceed humanity.

### **Summary B**

Mustafa Suleyman argues that today's surge in AI spending reflects a transformative technological wave rather than a speculative bubble. He notes that increased computing power has driven rapid capability gains, with AI already rivalling human performance in coding and poised to automate much white-collar work within 12–18 months. At Microsoft, his goal is AI self-sufficiency and the development of "superintelligence," defined as coordinated, autonomous systems that remain aligned with human interests. He stresses that safety and human control must take priority over speed in the AI race. Suleyman foresees a diverse ecosystem of specialized AI agents, major advances in medical diagnostics, and continued governance challenges as AI systems grow more capable and widely deployed.