Can AI make cutlery? – speech by Andrew Bailey

Given at the 389th Cutlers' Feast, Cutlers' Hall, Sheffield

Published on 21 May 2026

In his speech Governor Andrew Bailey highlights how UK economic growth has slowed since the Financial Crisis due to its legacy, bad shocks like Covid and wars, and the growing effect of declining population growth. He raises how economic growth can happen with AI likely the next "general-purpose technology". Though past experience suggests these gains may take time, depend on complementary innovation, and bring uncertain effects on jobs, making investment in skills essential.

Speech

It is my great pleasure to be at the 389th Cutler's Feast. It is exactly 30 years since Eddie George made the last Governor's speech at the Feast[1]. Master, the Company had a 70 year start on the Bank of England. We are a mere youngster.

From your founding in 1624, the company was the regulator of the cutlery trades in Hallamshire. The growth of the trade led to disputes over flouting of the company's rules, including going into partnership with foreigners. Was this people outside Yorkshire?

Senior Warden, you talked about the need for economic growth. I am going to join you on that subject. The title of my speech this evening is "Can AI make Cutlery"?

Let me start with a few key facts on growth.

For around fifteen years before the Financial Crisis (1990-2006), the average annual rate of growth of potential supply in the UK economy was around 2.8%. Productivity contributed 2.4pp and labour supply 0.4pp. For the roughly fifteen years since the Financial Crisis, the potential growth rate has fallen to 1.3% pa, of which productivity contributed 0.4pp and labour supply 0.8pp. The comparable figures for actual growth – the growth of demand – are 2.4% and 1.3% .

This has had a direct effect on living standards. In the UK, Between 1990 and the Financial Crisis, the average annual growth rate of national income per head was 2.0%, whereas post-Crisis it has been 0.6%.

What are the longer-run causes of growth in an economy? The list is quite long: physical capital accumulation; technological progress; human capital, including education and skills; trade; financial development; markets and competition; infrastructure; and institutional development, both public and private.

Theories of growth are often associated with two famous economists of the past, Adam Smith and Joseph Schumpeter. The terms Smithian and Schumpeterian growth are used, with a focus on trade and technology respectively.

Adam Smith used the pin factory to show that specialisation in production (the division of labour) leads to higher levels of productivity and hence growth. This was shaped by the size of the market served – the bigger the better. Trade expands markets, which enables productivity and the economy to grow by more.

Joseph Schumpeter wrote during the first half of the 20th century. He is known as the father of creative destruction, namely that capitalism is defined by waves of innovation, where new innovations and technologies replace older ones, and create new sources of growth in doing so[2].

Smithian and Schumpeterian growth are sometimes set up as competing explanations. This is not at all the case. They are complements[3].

Before coming to an example of innovation, let me say something about institutions, as a cause of economic growth. Institutions are important for economies and for growth, for instance because they define and operate property rights. But, they differ a lot across economies. And this, is where the Cutler's Company comes into the story.

One of the very big questions in the story of economic growth is why did Europe, and particularly Britain, industrialise first – how come we won the race? One view is that European institutions developed in a way that broke down family and clan based societies (something the authors argue did not happen for a long time in China). There was an important contribution here from trade organisations and corporations, like the Cutler's Company. It was from this institutional development that laws evolved, political stability was enhanced, and economic innovation and trade flourished. These overlapping social networks helped the sharing of ideas and knowledge. Master, today the Cutler's Company is still doing this[4].

Creative destruction involves innovation and obsolescence. Let me give an example close to the history of Sheffield and the Cutler's Company. Henry Bessemer's convertor made cheap molten steel a reality. Schumpeter had a flourishing turn of phrase. He described Bessemer as a real genius of the entrepreneurial kind, with a vision of the possibilities for cheap steel.

Bessemer came to metallurgy, which by his own account at the start he did not understand at all, due to his interest in developing artillery shells.[5] His original experiment succeeded, but using a non-phosphoric iron ore. Subsequently, he had to solve why it didn't work with phosphoric ones. Then, he had to confront the reality that hardly anyone wanted to buy a licence from him to use the idea.

On this point, Schumpeter raised himself to full rhetorical flourish. I will quote:

"Bessemer now took the line, much more consciously and according to reasoned plan than entrepreneurs do as a rule, of going straight into the citadel of the enemy, to Sheffield, in order to produce and undersell.

There, steel that was good, though not as yet very cheap (and here Schumpeter adds a footnote that it was not very cheap especially for other producers who had to pay for the licence from Bessemer, and it wasn't very cheap anyway because his productive apparatus was not the ideal of rationality), but eventually it saw the light of the market (from 1858) and was a financial success from the first"[6].

So, Master, here we are in the citadel of the enemy. I don't know how the Cutlers felt about the arrival of Henry Bessemer, perhaps not very joyous at first. But creative destruction had its way.

Now, I don't want to rain on Bessemer's parade but. It was a very important technological advance because the product – steel – was used across the economy. But it was not what is now known as General Purpose Technology, or GPT for short. A GPT is a technology (as distinct from a product like steel) with broad, economy-wide applicability that improves over time and enables complementary innovations across many sectors of the economy. To be clear, this use of GPT is not the same as Chat-GPT, where GPT stands for Generative Pre-trained Transformer.

GPTs have driven the cycles of technology-led growth, the heart of creative destruction. How many such cycles have we had? I would say the following qualify: the steam engine; the internal combustion engine; chemistry and the transformation of substances; electricity; and information processing and the internet.

I am going to go back now to the question, what has caused the slowdown in growth over the last fifteen years or so? There are quite a few candidates, and they are not mutually exclusive. Three possibilities are: the legacy of the financial crisis; the coincidence of bad shocks (Covid, the war in Ukraine caused by Russia; the latest outbreak in the Gulf); and the growing effect of declining population growth.

All of these are candidates, and may have a role. But, we should go back to Schumpeter , and to technology as a cause of growth. Three questions are relevant here. What is it about GPTs that causes growth cycles to end? Are there gaps between the cycles? And, is this a reasonable explanation of what has gone over the last fifteen years or so?

What causes GPTs to end? We are talking here about their contribution to growth, not their prevalence. The internet is everywhere, but as it has become so, it is natural that its contribution to growth will diminish. You can only get so much of it – honestly.

Are there gaps between cycles? Yes. We have seen this in the past, and they have been perceived as challenging periods. Economic historians have debated whether late Victorian Britain failed because other countries caught up with faster growth. I agree with those who say that it didn't fail, but it did need to reinvent itself[7].

One reason for such gaps is that it takes time to cook up the next technology.

Is this a reasonable explanation of the last fifteen years or so? Yes, I would argue. One clue is the declining contribution of productivity growth.

The obvious next question is what is the next General Purpose Technology? It's most likely to be Artificial Intelligence, though I would add AI operating with robotics.

This brings me to what may be the hardest question of the evening – how do we define AI? Bear with me – I do not claim to be an expert in this field[8].

As humans we are intelligent to the extent that our actions are expected to achieve our objectives. Machines are assumed not to have objectives of their own – they are a means of us furthering our objectives.

But we are often uncertain how to achieve our objectives. The critical feature of the human brain is that it can adapt to uncertainty. This is inductive thinking. It is distinct from deductive thinking, where we follow a pre-set rule. Where does AI fit into this? One step forward is to create machines that are able to process vastly more information more quickly than we can (eg process the whole of the internet). This is not doing anything we cannot do in principle, but in practice it is doing things we cannot do. But this is still fixed learning – fixed in the sense of the stock of information.

As humans we learn from experience – it's part of our inductive processing. This is where so-called Reinforcement Learning comes into AI. Here the machine learns through experience using trial and error. It takes actions, receives feedback, and adjusts its behaviour based on what worked and what did not. It is still applying general principles, but in doing so it is applying judgement based on its own learning. This creates a form of tacit knowledge. It applies general principles, while it also incorporates knowledge and judgement from its own

learning. The model is thus growing. You can see now that this is where the questions start to arise – is it really still set up to achieve our objectives or is it moving on itself and, in ways that may or may not be consistent with our intentions? That not a question for this evening - sorry!

But, what is relevant is how it operationalises all of this. Let's imagine we are Henry Bessemer. He had a finding – the initial experiment in London. He had to think through the constraints and opportunities of using it. Then, construct a plan to put it to work. Then he had to validate the steps in that plan (if only because his annoying banker may have asked for this). And then, he could get going and put the plan to work.

AI with Reinforcement Learning could get us much further through that sequence of steps with much less time spent, cost and human intervention, hence the productivity gains and its would-be status as a particularly potent GPT.

I can no longer resist asking the question in the title of this speech. Can AI make Cutlery? I did not take the question literally because I am of the generation that grew up watching Blue Peter in its early days. You may remember one of their familiar comments – "Don't try this at home".

Instead, I asked four different AI programmes the question – "Can AI make Cutlery"? They came back with basically the same answer[9]. Here's one; which is representative:

"Yes, AI can help make cutlery, but usually as part of a manufacturing system, not by itself. It can assist with design, automate production, inspect quality, and optimise packing and maintenance."

So, AI can improve the productivity of cutlery production, and is probably already doing so.

This brings me to the last part of my remarks. When are we going to get the growth from AI? I am going to take this in two parts. First, what are the characteristics of GPT innovation which can guide us to an answer? And, second, what can go wrong, - well, I'm a central banker, so we have to do this bit.

GPTs are enabling technologies – they spread across economies and make other things happen. The first invention often triggers a sequence of secondary innovations which spread its reach[10].

But, in most GPTs to date it has taken time for the main effects to come through. It took eight decades from James Watt's first steam engine to the major growth contribution of steam power showing up in the mid-nineteenth century. Likewise, while Thomas Edison opened his first power station in New York in 1882, it was a further four decades before there was a

meaningful boost to productivity growth. More recently, the economist Robert Solow made a famous comment in 1987 that: “You can see the computer age everywhere but in the productivity statistics” which became known as the Solow Paradox[11].

But, this may not be such a paradox, for several reasons. First, the initial invention typically doesn’t get there on its own – complementary changes need to come along and this takes time. Second, it is arguable that productivity may slow down in the early phases, because it is the experimenting phase. We may well be there on AI – I certainly am. Third, there may well be more measurement error in the early phases. We are not very good at measuring new things – recently the national statisticians have been working on how to value and measure data – which have a utility that has grown rapidly. And, we have to calculate to what extent GPT innovation increases the obsolescence of the old stuff that we used and how we should measure that.

One more point on why change can take time. An initial intervention which is not followed immediately by complementary further innovation can actually lead to the expansion of the old ways of doing things. In Sheffield, the early stages of nineteenth innovation actually led to the growth of the traditional craft-based domestic industry.

We can try to map all of this into what might happen with AI. The invention – innovation cycle may well have shortened, and maybe drastically so. Today, the life expectancy of a new model release is around three to four months. But, even so, that does not tell us how rapidly AI will shift the productivity numbers. AI may well ride to the rescue, but how quickly it will arrive in the productivity numbers is an open question. Really meaningful gains to productivity tend to come from wholly new products and activities, not so much from automating our existing tasks, and that tends to require further complementary innovations. There is a saying which is probably misattributed to the science fiction writer Arthur C. Clarke, that innovation is overrated in the short run and underrated in the long run. With AI, I think the innovation cycle is much shorter, but it may well still be somewhat longer before it contributes substantially to growth.

So, what can possibly go wrong? Schumpeter’s view of industrialisation and living standards was, rather simply, that in the end everyone came out ahead. Niccolo Machiavelli never exactly wrote that the end justifies the means, but Schumpeter’s argument heads in that direction[12].

A major issue here is what GPT innovation does for employment? The answer from history is not clear-cut. It depends on the form of the innovation and its impact. On one view, there are four main channels for the impact on employment[13].

- A displacement effect, where AI will replace people in tasks it can do more cheaply or efficiently. Chatbots replacing customer service agents would be an example (whatever

you think of them).

- A reinstatement effect – it creates new tasks where people are still needed. An example of this would be data scientists.
- A productivity effect – AI makes firms more productive, costs fall, output expands, and as the automation is partial, this increases the demand for labour. As with the reinstatement effect, the productivity effect expands the use of labour. An example of this would be AI helping doctors to analyse scans faster, more patients get treated and more doctors are needed.
- AI mostly replaces labour, and the share of profits increases relative to wages, employment can fall, but wages for certain skills can rise. This can increase demand for some highly skilled workers who manage AI systems.

The outcome in terms of employment therefore depends on which of these forces dominates, and the sequence by which they take effect.

History isn't a great help here. The first Industrial Revolution is usually dated from around the 1780s. While productivity growth did increase, real wages stagnated until the 1840s.

In this period, the wages of traditional skilled labour fell, and there was an increase in unskilled labour in factories, including child labour. The Luddites – machine destroyers - were a reaction to that, and were active at times in this area.

This period is also known as the Engels Pause. Friedrich Engels was closely associated with Karl Marx, and it was from this that Marx formulated much of his view on the dynamics of industrialisation and the eventual dictatorship of the proletariat. He was wrong, but it took time for that to become clear[14].

In contrast, mass factory production following electrification led to a more rapid increase in the demand for skilled labour and real wages and a fall in income inequality. The pattern of growth and income distribution was different.

The ICT and internet cycle coincided with a period of relatively low unemployment by historical standards. In the UK, the wage share of national income was relatively constant, while in the US it declined. The mix of skills in demand changed, and there was a growth of industrial production more broadly around the world. Offshoring rose, and there was strong growth in productivity in most parts of the world.

I would pull a few key points out here.

First, history shows that there is no single story on the relationship between major technology changes, economic growth, employment and living standards. However, shifts in the structure of the economy and employment are inherent to a technology-driven system of economic growth.

Second, the four part framework is a good way to think about the employment effects of AI. The answer is not pre-determined and it depends on choices we make – there is a dynamic element to the spread of GPTs. It is though a critical issue in the future of our economy. The impact of AI may well also be on different types of jobs from previous waves of GPT change. It may well impact more skilled jobs, and so-called entry level jobs.

I will end with one point – which Senior Warden goes back to your comments on the role of the Cutlers Company in supporting education. If there is one thing we can do which unambiguously will help on the employment question it is to invest in education and skills in AI.

Master, we live in interesting times. You have ensured that the Company is properly focused on the challenges of today, so I think you can look forward to another 389 Feasts.

Thank you.

I would like to thank Sarah Batten, Sarah Breeden, Karen Jude, Georgis Kyriaopoulous, Ben Martin, Dawn Plummer, Martin Seneca and Sam Woods for their comments and help in the preparation of these remarks.

---

1. Although Charles Bean spoke as Deputy Governor in 2009.
2. Schumpeter's best known work (amongst his many) is Capitalism, Socialism and Democracy, Harper and Brothers 1942. For a fuller account of his life and work, Thomas K. McGraw, Prophet of Innovation, Joseph Schumpeter and Creative Destruction, Harvard University Press, 2007. There is a very large modern literature on creative destruction, much of it based round the work of Philippe Aghion and Peter Howitt. Philippe Aghion and Peter Howitt, The Economics of Growth, MIT Press 2009. Philippe Aghion, Celine Antonin and Simon Bunel, The Power of Creative Destruction, Economic Upheaval and the Wealth of Nations, Belknap, 2023. Ufuk Akcigit and John Van Reenen, The Economics of Creative Destruction, New Research on Themes from Aghion and Howitt, Harvard University Press, 2023.
3. There can be dynamic welfare gains from increases in the rate of technological innovation and, without such gains from innovation, trade driven (Smithian) growth would encounter diminishing returns. Equally innovation in a closed economy would be less effective in terms of growth.
4. Avner Greif, Joel Mokyr and Guido Tabellini, Two Paths to Prosperity: Culture and Institutions in Europe and China, 1000-2000, Princeton University Press 2025. This view may seem to sit awkwardly alongside the view attributed to Adam Smith that industry groups tended to promote restrictive practices and restrict trade. But Smith in his later writings did recognise that in commercial countries like England, families did not stay together in the way associated with pastoral countries, and this was associated with commerce. Robin Douglass, A Moral Philosophy for Commercial Society, p127-8. In Paul Sagar (ed), Interpreting Adam Smith, Critical Essays, Cambridge University Press 2023.

5. Ulrich Wengenroth, Enterprise and Technology. The German and British Steel Industries, 1865-1895. Cambridge University Press, 1994, P17.
6. Joseph Schumpeter, Business Cycles, Volume 1, McGraw Hill, 1939, p372-3.
7. Joel Mokyr, Are We Living in the Middle of an Industrial Revolution, Federal Reserve Board of Kansas City Economic Review, Second Quarter 1997.
8. There is a very large literature on this subject. Two sources that I find useful, and suitably non-technical are: Stuart Russell, Human Compatible, AI and the Problem of Control, Penguin 2019; and Sebastian Malleby, The Infinity Machine, Demis Hassabis, Deep Mind, and the Quest for Super Intelligence, Allen Lane 2026.
9. The four were Chat GPT, Claude, CoPilot and Perplexity.
10. On the economics of GPTs, see: Elhanan Helpman (ed), General Purpose Technologies and Economic Growth, MIT Press, 1998.
11. Robert M Solow, We'd Better Watch Out, New York Times Book Review, July 12, 1987, p36.
12. This point is covered more extensively in, Carl Benedikt Frey, How Progress Ends: Technology Innovation and the Fate of Nations, Princeton University Press, 2025.
13. Daron Acemoglu and Pascual Restrepo, Artificial Intelligence, Automation and Work. In The Economics of Artificial Intelligence: An Agenda (ed) Ajay Agrawal, Joshua Gans and Avi Goldfarb, University of Chicago Press, 2019, P197-236.
14. The Engels Pause is described most recently in Carl Benedikt Frey, The Technology Trap, Capital Labor and Power in the Age of Automation, Princeton University Press 2019. The term originated in Robert C. Allen, Engel's Pause Technical Change, Capital Accumulation and Inequality in the British Industrial Revolution, Explorations in Economic History 46, no 4, 2009, p418-435. On the machinery question and Luddism: Maxine Berg, The machinery question and the making of political economy, 1815-1848. Cambridge University Press, 1980.

Andrew Bailey

Governor, Bank of England