

AUTUMN 2026

FINANCE

Presented by



EURO EXIM BANK
Facilitating Global Trade

MICHAEL BARR
DISCUSSES POLICY
CHOICES SHAPING
ARTIFICIAL INTELLIGENCE

PATRICK MINFORD
CONSIDERS IF AI WILL
TRANSFORM THE ECONOMY
AND UPEND OUR LIVES

ANDREW BAILEY
REFLECTS ON
INDEPENDENT CENTRAL
BANKS

21ST CENTURY FINANCE

Foreword

W

elcome to the Autumn edition of **FINANCE21**. This publication has been prepared in response to readership demand for an overview of the financial sector in these turbulent and unique times.

All aspects of the sector are examined, with the most respected authors providing the reader with the most comprehensive information available. Our brief is to provide all the data necessary for the readership to make their own informed decisions. All editorials are independent, and content is unaffected by advertising or other commercial considerations. Authors are not endorsing any commercial or other content within the publication. ■

CONTENTS

The institutional form of independent central banks

Andrew Bailey argues that independent central banks rely on more than technical expertise: they have a long tradition of accountability. Central bankers must explain their decisions or risk appearing to be an unrepresentative elite, especially amid the challenge of populism

AI takeover or slow grind?

Will AI suddenly transform the economy and upend our lives? Or is the prospect more mundane? Patrick Minford considers the effects of earlier systemic innovations and believes AI is unlikely to be different from other technology revolutions

AI and financial inclusion

AI has the potential to transform lives and the US economy. Michael Barr warns that artificial intelligence could dramatically narrow or widen economic inequality in the US, depending on the policy choices society makes in the coming years

How the rise of AI matters for fiscal policy

Predictions of the effects of AI on the economy and society tend to focus on productivity growth, job losses, and shifts in the distribution of income. Karen Dynan, Douglas Elmendorf and Louise Sheiner examine various potential scenarios for the macroeconomic effects of AI on US fiscal policy

CONTENTS

A new age of capital: growth, sovereignty and AI

AI is the project that could drive European capital markets forward to channel savings where they are needed, says Christine Lagarde. It needs capital on a scale that only markets can provide, and the stakes for Europe could hardly be higher

Supercharging AI into explosive growth

If frontier AI models can be used to develop their next models, this may lead to an 'intelligence explosion' and give a large boost to technological advancement. Davidson *et al* develop a framework for understanding how AI may transform the economy through new feedback loops

Data centres are laying the next financial mine

Yash Kalash argues that the AI infrastructure boom is creating financial transmission mechanisms reminiscent of past crises—not because AI will fail, but because its financing may not endure long enough to repay the infrastructure built around it

Some questions that a BoE review should ask

As the Bank of England nears 30 years of operational independence, a broad review is likely. Sushil Wadhwani proposes five measures to reduce future quantitative-easing losses and, as markets still doubt the Bank will meet its inflation target two years after the Bernanke Review, steps to strengthen credibility

CONTENTS

From vision to delivery

The ECB wants central bank money to serve as the monetary anchor for tokenised markets. Piero Cipollone calls for common standards, central bank money and coordinated regulation to prevent blockchain-based finance from becoming a collection of disconnected platforms

Achieving a European Banking Union without increasing risks

European banking remains fragmented, leaving billions of euros inefficiently allocated. Gyöngyi Lóránth, Anatoli Segura and Jing Zeng argue that a common deposit insurance system can encourage more risk-taking, but this should be accompanied by sufficiently strict supervision

Reforming EU banking regulations

EU banking rules are under pressure for simplification. Jesper Berg and Hans Geeroms argue that reforms should cut complexity without weakening financial safeguards

Modernizing financial regulation

The financial system is constantly evolving, and the regulatory and supervisory framework must keep pace. Michelle Bowman discusses the work of the US and urges global financial watchdogs to be flexible

CONTENTS

A strategic lever to boost global decarbonisation

The Paris Agreement provides a framework for international carbon trading. Georg Zachmann and Annabelle Weisser and argue that international carbon credits can support the EU's 2040 target if buyers' club rules and leverage ensure real mitigation gains abroad

Securing European prosperity

In the State of the Union address, Ursula von der Leyen shared the key priorities and flagship projects for the year to come, building on the progress the EU has made in recent years and setting out the priorities for the year ahead and her vision for a stronger, more independent and resilient European Union

Rebuilding the world order

As the postwar international order fragments and universal agreement becomes increasingly elusive, policymakers face a pressing question: what form of cooperation can replace the postwar hegemon-led, rules-based system?

Dennis Snower examines two complementary answers

Redrawing the global trade map

Liberation Day raised an immediate question in policy circles: would other countries follow suit? Barthélémy Bonadio, Andrei Levchenko and Nitya Pandalai-Nayar study the implications of the US withdrawal from free trade

CONTENTS

Europe's corporate tax trilemma

The EU's implementation of global anti-tax avoidance rules and the global minimum tax is widely regarded as a success. Madalena Barata da Rocha, Roel Dom, Pascal Saint-Amans and Bo Sangers consider the trilemma between tax sovereignty, revenue protection and investment neutrality that has emerged

Europe's federated democratic learning state

Europe is moving towards invoice-level electronic VAT reporting. Thiemo Fetzer argues that, with AI-assisted processing of the information contained in electronic VAT invoices, the tax system can become not just a compliance and revenue-raising instrument, but a protected intelligence map of the economy



www.euroeximbank.com
EURO EXIM BANK
Facilitating Global Trade

www.euroeximbank.com

HELPING YOU STAY AHEAD OF THE GAME

Letter of Credit | Standby Letters of Credit |
Bank Guarantees | Proof of Funds



We aim to accelerate your import &
export business transactions and
connect you to new opportunities

Representative Office (UK), Euro Exim Bank Ltd, Suite 1C, 1 Oaks Court,
1 Warwick Rd, Borehamwood, Hertfordshire, WD6 1GS, United Kingdom.

Registered Office (WI) Euro Exim Bank Ltd, Office 2, 1st Floor, La Place
Creole Building, Rodney Bay, Gros Islet, St. Lucia, LC01 101, West Indies



info@euroeximbank.com



+44 208 207 2868 (UK)

+1 758 450 8349 (WI)

Euro Exim Bank Ltd (Representative Office) is incorporated in the UK under company registration no. 07520196, with the registered address at Suite 1C, 1 Oaks Court, 1 Warwick Rd, Borehamwood, Hertfordshire, WD6 1GS, United Kingdom. This is the UK operation of Euro Exim Bank Ltd (Registered Head Office), Office 2, 1st Floor, La Place Creole Building, Rodney Bay, Gros Islet, St Lucia, WI.

The institutional form of independent central banks



Andrew Bailey argues that independent central banks rely on more than technical expertise: they have a long tradition of accountability. Central bankers must explain their decisions or risk appearing to be an unrepresentative elite, especially amid the challenge of populism

would like to reflect on independent central banks as institutions, and how their authority is grounded and is challenged. On 3rd September 1780 (246 years ago) a 23-year-old aide to George Washington wrote that: *“the Bank of England unites public authority and faith with private credit and hence we see what a vast fabric of paper credit is raised on a visionary basis.”* That 23 year old was Alexander Hamilton¹.

The meaning of central bank independence is not just a technical point about monetary policy or financial stability, or just a point about the design of public institutions. It is about the relationship between societies and the institutions they create to serve them. This brings to mind Douglass North’s definition of institutions as the rules of the game in a society or, more formally, the humanly devised constraints that shape human interaction².

Institutions are not just organisations. They are durable arrangements through which societies solve collective problems. To do this, they must embody understandings about how authority should be exercised, how conflicts should be managed, and how common interests should be protected.

Independent central banks are among the most distinctive institutions modern societies have created. They exercise important powers. Their decisions influence borrowing costs, saving decisions, employment prospects, investment plans and the broad stability of the financial system. Their powers, and how they are used, are always the subject of debate, and that’s fair enough given the importance.

Now, there are critics of these arrangements. Some critics argue that central banks disproportionately reflect the interests of financial markets and financial institutions. In a similar, but distinct, spirit others argue that central banks are too detached from the preferences of citizens. These criticisms often differ in their origins and motivation, but they converge on a common question. Why should an unelected institution possess such significant authority?

My purpose is not just to defend central bank independence, though you can be assured I am happy to do that robustly. Rather, it is to reflect on the deeper intellectual traditions from which it has emerged and to consider what those traditions can tell us about the challenges we face.

When economists explain central bank independence, they often begin with questions about the credibility of policy regimes and their time consistency. These arguments are important. They explain why in many countries independence emerged because governments at crucial times had demonstrated incentives to pursue short-term monetary policies that ultimately undermined economic stability, and to undercook financial stability outside more immediate proximity to a financial crisis.

Central bank independence does not mean detachment from democracy. It means insulation from short-term political pressures within a democratic framework

But they do not fully answer a more fundamental question: why should society trust an institution that operates with a degree of independence from day-to-day political control? To answer that question, it is useful to look further back.

One important tradition begins with the French philosopher Montesquieu³. Writing in the eighteenth century, Montesquieu was preoccupied with the problem of concentrated power. His great concern was how societies could prevent authority (in that case royal authority) from becoming arbitrary. His answer was that power should be moderated through institutions that stood between rulers and citizens.

Parliaments, courts of law and other intermediary bodies performed this role. They constrained power, not simply through legal rules but because they possessed recognised standing within society. They had identities, traditions and sources of legitimacy independent of those they were intended to restrain. For Montesquieu, these roots mattered. Institutions could not simply be invented and expected to command authority. They required legitimacy that grew over time.

At first sight, modern central banks appear rather different. A central bank derives its authority primarily from legislation. Its independence is granted through statute. It does not rest upon inherited privilege, social rank, or customary authority. Modern central banks are deliberately designed to be technocratic institutions.

And yet, the questions Montesquieu asked remain relevant. His central concern was how societies create institutions capable of constraining the exercise of power for the common good. In that sense, the challenge he identified remains with us.

A second intellectual tradition is also relevant. This tradition emerged from English political thought in the seventeenth and eighteenth centuries, and from debates about sovereignty, money and commerce. As David Hume put it, in the seventeenth century commerce, for the first time, became a *“reason of state.”*⁴

Thomas Hobbes provided one starting point in this development⁵. In the world of Hobbes, the authority to determine the value of money belonged to the sovereign. Just as the sovereign commanded armies and administered justice, so too the sovereign controlled the monetary system of the day, such as it was. Given this overriding concern with political order and stability, for Hobbes dividing authority over money risked dividing authority itself.

But this view was challenged. John Locke argued against the idea that governments should be able to alter the value of money according to political convenience⁶. During the monetary controversies of the 1690s, Locke insisted that arbitrary changes in the value of money damaged trust, disrupted commerce and undermined the public faith upon which economic exchange depended.

This marked a significant intellectual shift. Money was no longer seen just as an instrument of sovereign authority. It became something that served a broader public purpose. Its stability mattered not merely because government or the Sovereign desired it, but because society required it.

David Hume developed this argument further⁷. Hume observed that commerce had become central to the prosperity and strength of nations. Economic life was no longer peripheral to statecraft. It had become one of its defining concerns. Stable money and reliable credit systems became essential foundations for expanding commerce. They enabled people to transact, invest, save and plan with confidence.

In this tradition, monetary stability became a public good. It is from this intellectual lineage that modern central banking most directly emerged. The founding of the Bank of England in 1694 illustrated this evolution, reflecting two powerful forces⁸.

The first was the rapid expansion of trade and commerce. As markets became more sophisticated and more interconnected, the need for a stable and reliable monetary and financial framework grew.

The second was the necessity of financing war. In 1694, England's conflict with France created urgent demands for finance. The arrangement that established the Bank of England enabled private capital to be mobilised in support of public objectives.

One of the interesting points in *The Spirit of the Laws* is that Montesquieu linked England's financial credibility to its constitutional arrangements (comparing it to France). In England, public credit was not merely a matter of financial technique; it depended on institutions capable of generating public trust. This interpretation is entirely consistent with the quote from Alexander Hamilton.

But the opposite view is also powerful, as set out by John Pocock and summarised by Istvan Hont:

"The understanding of commerce and commercial society in late seventeenth century England was most advanced by just those thinkers who saw the political health of the English nation as jeopardised by the pervasive corruption that resulted from the intrusion of the demands of the war economy into the sphere of domestic politics."

There was a growing awareness of the distractions which the new of kind of war finance, the running of a huge public debt, might cause to the balance of the English constitution. A political and military regime financed by

the public debt implied a danger of the rising in influence of a financier class, which could exercise power without enjoying a proper and legitimate place in the politics of the country⁹.

One way to describe the importance of the growth of commercial societies is to show how they led to the emergence of some of the most profound issues of political economy, namely:

- The rise of nation states as economic and commercial agents in their own right;
- The need for a properly constituted system of sovereign power to settle internal and external disputes;
- The demand for state finance;
- The role of justice in economic distribution;
- The growth of elaborate divisions among economic interests;
- Greater complexity and opacity in the workings of economies; and
- The development of tools for authorities to modify the workings of markets, and the risk of a variety of motives for doing so¹⁰.

These issues remain at the heart of economic policymaking today. The early Bank of England was an important part of the development of commerce. It is however important to recognise that the original Bank of England looked very different from the institution we know today.

It was not conceived as an independent guardian of monetary stability. It emerged from a partnership between government and a relatively small and concentrated group of financially sophisticated investors. And, unlike the modern institution, the early Bank was not removed from politics¹¹.

The origins of the Bank align as much with Mancur Olson's later analysis of organised interest groups as with any modern idea of an independent public institution¹². My point is not to criticise those origins. They reflected the realities of the time. But they do remind us that central banking did not emerge fully formed. It evolved gradually as understandings of money, credit, commerce and public authority evolved. Nonetheless, the big themes we wrestle with today were beginning to emerge at that time.

The modern central bank is therefore the product of a long historical process. On one hand lies the Montesquieu tradition, concerned with limiting arbitrary power through institutional checks. On the other lies the tradition of Locke and Hume, concerned with establishing money and credit as public goods upon which society depends.

Modern central banks bring these traditions together. Their essential purpose is to safeguard the stability of money including the resilience of the financial system. These are not interests that belong exclusively to any group. They are public goods from which society as a whole benefits.

The important distinction between the public interest and the public good, as set out by Brian Barry, is important in this respect¹³. Individual citizens naturally have different preferences, priorities, and interests. Democracies exist in part to sort out those competing claims. But some goods are valuable to the public as a whole precisely because they are universally shared.

Monetary and financial stability is one such good. Whatever our views on taxation, spending, regulation or public services, we all operate within a common monetary and financial framework. The value of money is something upon which every household and every business rely. To be effective in the sense of trust and legitimacy in society, support for such a public good should come from the plurality of society rather than any single member group.

This does not mean that monetary policy and financial stability are devoid of distributional effects. Interest rates affect borrowers and savers differently. Financial stability policies can affect sectors and institutions in different ways. Public policy inevitably has consequences that are not distributed uniformly.

But recognising those consequences should not obscure the broader purpose that these policies serve. The objective is not to favour one group over another. It is to preserve the monetary and financial conditions that make sustained prosperity possible.

In that sense, central bank independence represents a modern adaptation of Montesquieu's thinking. The institution exists not to accumulate or counter arbitrary power but to protect the value of money from short-term pressures that could damage the public interest. But this arrangement also creates an enduring challenge.

Unlike the institutions that Montesquieu had in mind, central banks do not possess deep social roots of their own. Indeed, they must be careful not to become identified with any single interest group. A central bank that became closely aligned with particular sectors or political movements would undermine the very rationale for its independence, and therefore its legitimacy.

As John Dunn has argued, this gets us to the issue of the constitutional theory of how power can be institutionalised and thus be implemented effectively in practice. This is the critical relationship between *“the*

constitutional order of power and the pragmatic determination of economic policy.” It’s the modern version of the problem Montesquieu identified and sought to solve.

As Dunn notes, constitutional order can appear very settled and continuous, whereas economic policy can appear anything but at times. Writing in the mid-1980s, Dunn argued that a constitutional order which rendered the benign conduct of economic policy more likely than not had not yet been constructed. Solving this challenge has been critical for the legitimacy of modern central banks¹⁴.

So where does legitimacy come from? The answer, I believe, is now straightforward. Modern independent central banks are not independent of democratic government in the important sense that their authority originates in legislation or the like enacted by elected representatives. In the Bank of England’s case, Parliament defines the objectives, establishes the powers, and retains the authority to amend the framework within which we operate. This role is critical.

This brings me to the contemporary challenge. Across many countries we have witnessed growing scepticism towards public institutions generally. In the case of central banks, as I mentioned earlier, some critics argue they are too close to financial interests and that they stand in the way of popular preferences.

This challenge is not new, but some aspects of it are sharply set out at present, particularly in the area of financial regulation. The argument goes that central bank policies affect the distribution of returns, wealth and welfare directly, which pulls economic interests back into the picture. This concern applies more to financial stability policy than to monetary policy, though monetary policy is not immune from it.

One version of the criticism is that the central bank remains too close to financial interests. A central bank's usual response is that the public good of stable money substantially outweighs any distributional effects. But this response can return the argument to Olson's thinking, since the financial world does contain well organised and concentrated interests.

The irony is that central banks can, and today do, face an opposite accusation from the financial interests themselves, who argue that regulation has gone too far. So, criticism arrives simultaneously from two directions that contradict each other.

The fact that criticism can come simultaneously from opposing directions does not prove that central banks are correct. On their own, two perceived wrongs don't make a right. But it does underline that independence often requires making decisions that do not align with the private interests of any constituency or group, and that at times resisting pressure from the concentrated interests that Mancur Olson identified is part of the task of protecting a public good like the value of money.

A second challenge comes from populist political movements. As Jan-Werner Müller has argued, populism often rests on the claim that a particular political movement alone represents the authentic will of the people¹⁵. Put thus, the public interest is something that a single ideology can capture and deliver directly. This effectively collapses Barry's distinction between the public interest and the public good.

Any institution seen to get in the way becomes an unrepresentative elite standing between the people and their will, and thus an obstacle to popular sovereignty. This is a serious challenge. We have developed systems of government (in the broadest sense of this term) in which legitimacy rests in the plurality of society, not in the preferences of any single group within it.

Courts of law, universities, regulatory bodies and central banks all derive authority through forms of democratic delegation set within a framework of that plural society¹⁶. Maintaining the trust that goes with that delegation matters deeply.

Central banks cannot take legitimacy for granted. They must explain their decisions clearly. They must engage openly with the public. They must remain accountable to elected representatives for their statutory objectives. In our case that means the Treasury Select Committee of the House of Commons. Through this engagement and accountability, we must demonstrate how our actions serve the public good, in its traditional meaning.

Let me conclude. Institutions that exercise significant authority should expect scrutiny. Such scrutiny is a sign of democratic health. But we should also recognise what is at stake.

Sustained economic prosperity depends upon confidence in the future. Households must be able to spend and save. Businesses must be able to produce and invest. Financial markets must be able to allocate capital efficiently. All of these activities rely upon confidence that the value of money will remain stable including that the financial system will remain resilient.

Those conditions cannot be created overnight. They depend upon institutions that are trusted, durable and credible. Central bank independence does not mean detachment from democracy. It means insulation from short-term political pressures within a democratic framework. Its legitimacy derives from a parliamentary delegation and accountability goes with that independence. It is the basis on which independence legitimately rests.

In our own time, as in previous generations, the challenge is therefore not to weaken that foundation but to strengthen it: through accountability, through transparency, through public engagement and through an unwavering commitment to the public good of sound money that independent central banks exist to serve. ■

Andrew Bailey is the Governor of the Bank of England

Endnotes

1. Alexander Hamilton to James Duane, 3 September 1780, in *The Papers of Alexander Hamilton*, ed. Harold C. Syrett, vol. 2, 1779–1781 (Columbia University Press, 1961), 400–418.
2. North, DC, *Institutions, Institutional Change and Economic Performance*, (Cambridge University Press, 1990).
3. Montesquieu, CL de S, *De l'esprit des lois* [The Spirit of the Laws], (Barrillot and Fils, 1748).
4. David Hume, "Of Civil Liberty" in, *Essays Moral, Political and Literary*, (Alexander Kincaid, 1741)
5. Hobbes, T, *Leviathan*, (Andrew Crooke, 1651).
6. Locke, J, *Some Considerations of the Consequences of the Lowering of Interest, and Raising the Value of Money*, (Awnsham and John Churchill, 1692),
Locke, J, *Further Considerations Concerning Raising the Value of Money*, (A and J Churchill, 1695).
7. Hume, D "Of Commerce", in *Political Discourses*, reprinted in *Essays, Moral, Political, and Literary*. (1752).
8. Dickson, PGM, *The Financial Revolution in England: A Study in the Development of Public Credit, 1688-1756*, (Macmillan, 1967).
9. Istvan Hont: "Free trade and the economic limits to national politics: neo-Machiavellian political economy reconsidered; in John Dunn (ed). *The economic limits to modern politics*. (Cambridge University Press 1990). P49-50. JGA Pocock, *The Machiavellian Moment: Political Thought and the Atlantic Republican Tradition* (Princeton University Press 1975).
10. John Dunn, *The economic limits of modern politics*. (Cambridge University Press, 1990), p37-40.
11. Lord Macaulay, observed that "the weight of the Bank, which was constantly balanced in the scale of the Whigs, almost counterbalanced the weight of the Church, which was as constantly in the scale of the Tories". Thomas Babington Macaulay, *The History of England from the Accession of James II*, vol IV, ch XX (various editions), p556.
By 1781, Prime Minister Lord North forcefully advocated for the Bank's charter renewal on the basis that the Bank had become so important to national credit and public finance that it could no longer be treated as an ordinary private

corporation and had assumed de facto constitutional status, or at least a distinctive status within the executive branch. North had started his career as a Whig, but by this time was more aligned with the Tories.

12. Olson, M, *The Logic of Collective Action: Public Goods and the Theory of Groups* (Harvard University Press, 1965).

13. Barry, B, *Political Argument*, (Routledge and Kegan Paul, 1965)

14. John Dunn, "Capitalism, socialism, and democracy: compatibilities and contradictions" In Dunn (ed), *The Economic limits to modern politics*. p195-219.

15. Müller, J-W (2016), *What Is Populism?*, (University of Pennsylvania Press, 2016).

16. On this point, see: Rawls, J, *A Theory of Justice*, (Belknap Press of Harvard University Press, 1971).

This article is based on a [speech](#) delivered at LSE TRIUM Anniversary Conference, September 4 2026.

AI takeover or slow grind?

Will AI suddenly transform the economy and upend our lives? Or is the prospect more mundane? Patrick Minford considers the effects of earlier systemic innovations and believes AI is unlikely to be different from other technology revolutions

There is regular comment in the press about the effects of AI on employment, productivity and growth, as well as concerns about general social effects. Opinions vary between fears of a growth explosion accompanying mass unemployment, with huge social tensions - a view that dominates the political debate - and at the opposite end reassurance that the progress of AI will be slow as it is gradually absorbed into the economy's structure.

At the same time there is general agreement on the need to regulate AI output for social reasons in a variety of ways; this includes its use by young people on smart phones and its use by bad actors, such as terrorists. The latter agreement essentially just extends existing law to the use of AI; the question that has yet to be settled is quite how this is to be done, whether by setting up new regulation or by letting existing laws and law courts mandate action.

This issue of social regulation is important but not my focus here. Rather it is what concerns us economically: the other controversy over the economic effects of the AI revolution. Here we can be helped by the history of employment and output in the face of earlier systemic innovations, dubbed General Purpose Technologies, or GPTs.

These include the invention of spinning and weaving machinery, the internal combustion engine, railways, electricity and the computer. These five GPTs arose between the 18th and the 21st centuries and it is instructive to see how trends in employment and output reacted to them in the UK economy, for which we have the longest data series on their impact. The charts that follow - all taken from the Bank of England's historical databank - tell the story.

What we see from all these trends is that productivity and employment followed fairly steady upward trends throughout the three centuries of massive technological changes due to these GPTs. Most tellingly, the five-year-moving average of TFP (overall productivity) growth – see Figure 2 - fluctuated since 1750 at a rate between 0 and 3%.

Figure 1. UK labour productivity from 1860



Source: Bank of England via FRED®

Figure 2. UK TFP growth since 1750





Figure 3. UK unemployment rate since 1750




Source: Bank of England via FRED®

Figure 4. UK Labour share of GDP from 1750



Source: Bank of England via FRED®

Figure 5. UK employment from 1750



Source: Bank of England via FRED®

There were many episodes of high unemployment but these are largely responses to financial and other macro shocks, 'cyclical' in other words, even if on occasions linked to 'bubbles' associated with GPT-related behaviour- such as the railways overbuilding in the middle of the nineteenth century (Figure 3 shows the sharp associated rise in 1850, which was soon over).

The biggest rise in unemployment was in the interwar period, and was related to monetary problems with the gold standard, and later the Great Depression - nothing to do with GPTs. Then the high unemployment of late 1970s was related to the oil crisis of 1975 and the battle against double digit inflation.

AI is unlikely to be different from other technology revolutions. It will take time for the world economy to assimilate the new methods into general business practice as new products and ways of working are found that use it effectively and safely

Hence the evidence suggests that previous GPTs did not cause employment dislocation but rather were absorbed into the economy in a rather steady way, gradually raising productivity and creating as many jobs as they displaced. This evidence therefore supports the 'slow and steady absorption' view of the AI revolution rather than the 'big explosion' view that is found in much of the political debate.

How can we account for this stability of these trends in the face of the massive technological shifts implied by these GPTs. It seems that the application of these GPTs required largescale changes in product types and methods of production to be fully profitable. Thus, for example it was not until Ford invented the Model T and scaled it to mass production that the full profitability of the internal combustion engine and electricity was achieved by business.

In the same vein Robert Solow famously joked in the *New York Times* in 1993 that the computer could be seen everywhere except in the productivity figures; yet now work across the whole economy relies on the computer, with massive cumulative productivity effects in total. This 'Solow paradox' has been seen to be resolved by the slowness with which new uses are found for a new GPT, as described above.

When we turn to AI, recent work (eg. [Demirer et al \(2026\)](#)) has found that AI is widely used by individuals in their work, and has become a huge cost to businesses as a result of the large use of 'tokens', but has not resulted in much increase in firms' output.

Demirer *et al* suggest that AI will become profitable as new businesses built entirely on AI processing take over from currently dominant firms. Hence AI is following the usual path of GPTs with productivity growth being gradually raised as it is absorbed by industry into new products and methods.

This impression is strengthened by the recent commentary from business consultants. Thus, the latest [Wharton report](#), which is more bullish than most, finds that only a third of businesses surveyed are making a 'significant' return on AI usage, this being led by smaller businesses. The latest McKinsey report (McKinsey Global surveys on the state of AI, 2017-2025) is sceptical of AI progress: over half of businesses surveyed are getting a revenue rise or cost reduction from AI, two thirds experimenting but not scaling up; no profit is reported yet.

The latest Cambridge Judge Business School [survey](#) of 600+ finance companies says only 40% report any profit and most only use AI for back-office tasks. [Accenture](#) says 90% of businesses are making no money so far. All this supports the 'slow and steady' view of the AI revolution.

One of the things that accompany early enthusiasm for a GPT is high investment by providers in the face of slow adoption. This creates conditions for a collapse in equity values of these early investors- usually dubbed the bursting of a 'bubble'. This happened with the railways GPT overbuilding of the mid nineteenth century and also the dotcom boom of the 1990s.

The current massive wave of investment by AI providers in computers, cloud provision and data centres, may well be leading up to a similar crisis of overprovision. In 2025 global corporate AI investment was \$582 billion, around 2% of US GDP where most of it is being commissioned, around half of it funded by venture capital as equity, the rest being borrowed.

This is still less proportionately than railways investment which reached 7% of UK GDP in 1847, but it may be enough to trigger a bust. In addition, the US Federal government spent \$3.3 trillion on AI R&D in 2025, about 10% of GDP, whereas the government was not involved with railways investment. This difference, together with the wide

popular backlash against the private investment in data centres because of their absorption of electricity and water resources, may be enough to avert a bubble and bust.

Looking at other details of how the AI revolution is being rolled out, it is clear enough that AI is easily used with strongly productive effect in certain areas such as repetitive search for new drug combinations, military target selection from wide selections fed in, and driverless cars.

These are areas where massive data troves can be fed in and these then can be used for tasks that search through these in a guided way. The problems encountered here are also daunting, mostly involving 'hallucination' where the AI picks a wrong combination as a solution. With drugs search this can be easily dealt with since the 'discovered drugs' can be checked by humans later.

However sometimes this is impossible. For example, in military targeting the whole point is to bypass the human element to achieve speed. Yet a wrong target such as the girls' school selected in the bombing of Iran cannot be avoided once selected; the costs of such an AI hallucination are potentially huge - in terms of reputational damage, not to speak of the direct cost in lives. Much the same applies to driverless cars; one mistaken turn resulting in a terrible accident, could put back the prospects for driverless cars for a generation.

These are just examples of the ways in which using AI can fail to be valuable until there are future developments in product uses that avoid such problems or at least are no worse than the effects of endemic human error.


It is probably for this sort of reason that the latest job reports in the US and elsewhere have not suggested a collapse in jobs for new labour market entrants as once feared - see Figure 6 for total US employment, which shows steady job creation apart from in the Covid episode.

The data for employment across the whole OECD tell a similar story - Figure 7. New entrants are needed to supervise the use of AI and avoid hallucinatory mistakes that can severely damage a firm's reputation. These entrants in turn need to be trained for this task by existing employees that can pass on the necessary skills and cautionary philosophy.

To conclude, AI is unlikely to be different from other technology revolutions. It will take time for the world economy to assimilate the new methods into general business practice as new products and ways of working are found that use it effectively and safely. To do this people will be needed in employment to fashion the new world and train the new employees in the new methods. Even with AI the world economy will not turn on a dime. ■


Patrick Minford is Professor of Applied Economics at Cardiff University

Figure 6. Total US nonfarm employment




Source: US Bureau of Labor Statistics via FRED®

Figure 7. Total OECD employment rate



Source: Organization for Economic Co-operation and Development via FRED®



Download World Commerce Review Autumn 2026, a fully interactive review of global trade and finance for the last quarter.

Download



**SCAN TO
DOWNLOAD**

AI and financial inclusion



AI has the potential to transform lives and the US economy. Michael Barr warns that artificial intelligence could dramatically narrow or widen economic inequality in the US, depending on the policy choices society makes in the coming years

Our focus is financial inclusion, and something that will likely have great consequences for financial inclusion and our economy more broadly in the years ahead is artificial intelligence (AI). As I have explored in a number of speeches over the past several years, AI has the potential to transform lives and the US economy, possibly empowering workers to be more productive, with lower- and middle-income workers benefiting the most¹.

But it is also the case that AI may instead exacerbate inequality, eliminating some lower- and middle-income jobs while boosting the income and wealth of higher-income individuals. Since we don't know which of these futures will come about, it is useful to use potential scenarios, as I've done previously with respect to AI and the economy.

Every major technological advance has had profound effects on labour markets and the economy. Many workers have suffered from these technological changes, while many other workers have seen new opportunities emerge. In the long run, technological advances tend to broadly raise living standards by creating more jobs than they destroy and increasing productivity.

But transitions and outcomes can vary, and in the period following the mass adoption of a general-purpose technology—such as electricity, the telephone, and internet-enabled personal computers—the number of people dislocated and the extent of the harm they may suffer can be large and persistent.

Balancing those scales of costs and benefits involves examining whether the benefits are broadly shared or concentrated. Past experience has shown that technological leaps forward can raise living standards. But when the benefits are concentrated among relatively few people, technology can widen inequalities of income and wealth, especially during the transition period.

Widespread adoption of the internet raised the productive capacity of our economy and broadly raised living standards, but it also likely exacerbated inequality because it benefited information-intensive jobs (such as accountants) more than other jobs (say, construction workers).

The policy challenge is therefore not simply to observe the development and deployment of AI, but also, as a society, to consider policies related to AI and its effects on education, job training and workforce development,

Every major technological advance has had profound effects on labour markets and the economy. Many workers have suffered from these technological changes, while many other workers have seen new opportunities emerge

competition, tax policy, and other areas that allow the gains from AI to be shared across workers, households, and communities, rather than accruing to a small group of firms and investors². Of course, these policies are not within the remit of the Federal Reserve but rather for other policymakers to consider and decide.

The question I would like to consider is whether AI will likely help narrow inequalities of income and wealth, supporting advances in financial inclusion, or widen those inequalities, undermining the recent gains in financial inclusion that we rightly celebrate today.

Understanding inequality

Let's start with understanding income and wealth inequality. To understand income inequality, it helps to break income into its components. The largest component is labour income. Disparities in labour income across individuals reflect the supply of and demand for their skills, their productivity, and their time spent at work—all things that will be affected by AI.

Income also includes earnings from capital and investments and, thus, includes the concentration of ownership in firms, such as AI companies. In 2024, the highest-earning one-fifth of US households earned 52 percent of all income, and the bottom 20 percent earned only 3 percent³. In 2024, the United States was the sixth most unequal of the countries in the G20⁴.

Wealth inequality is a function of the distribution of ownership of assets—land, goods, businesses, intellectual property, and other investment assets. The bottom one-half of US households hold less than 3 percent of wealth, the top one-tenth hold 59 percent, and the top one-tenth of a percent hold 15 percent⁵. When returns from investments are reinvested, wealth naturally compounds. As a result, those who already own appreciating assets

often see their wealth grow much faster than households that rely primarily on wages, widening the gap between the 'haves' and the 'have-nots'.

A central question is whether AI will expand opportunity by giving more people access to valuable skills and productive work, or whether it will reinforce advantages that are already concentrated among a smaller population. Inequality matters not just for workers today, but it is also closely connected with an important aspect of the American Dream—the expectation that in the future, our children will be able to make better lives for themselves, including through rising living standards⁶.

How AI could widen inequality

Let me start with the possible ways in which AI could widen inequality.

Automation and labour displacement

First, AI could lead to labour displacement. Something that is at the top of mind for most people, especially younger workers, is the concern that AI could drastically reduce the demand for them. AI might disproportionately affect new entrants to the labour market. According to a well-known paper by Claudia Golden and Larry Katz, in some previous technological waves, the benefits have tended to improve outcomes for more-skilled and more-educated workers⁷.

But in this scenario, AI could harm not only less-skilled workers, but also younger college-educated workers whose skills are more easily replicated by AI than in prior technological waves. Moreover, workers who use AI more intensely might gain the most, and workers who use the most advanced and expensive AI models might win out over those who use baseline models.

In the Federal Reserve's most recent Survey of Household Economics and Decisionmaking, 43 percent of workers with a graduate degree reported using AI in the previous month, compared with 10 percent of workers with a high school degree or less⁸. The survey further found that workers who used AI were more likely to say that it would improve their careers than replace their jobs.

Measures of exposure to generative AI also suggest that higher education and higher-paid workers are much more exposed to generative AI⁹. The consequences of these facts are not yet clear. If AI mostly substitutes for labour, then not being exposed to AI would be positive for such workers; however, if AI augments existing jobs, then workers not exposed to AI would be left behind.

There is substantial uncertainty about how the labour market will evolve. As of right now, there has been little evidence of economy-wide job displacement from AI¹⁰. Yet there is some evidence that AI may have made job entry harder for young workers in some job categories¹¹. Given both the history of major technological advances and how early we are in the timeline of AI adoption, it is important to consider the full range of possible future effects, including the potential for more widespread labour displacement.

Potential concentration

Another concern is concentration. A high degree of market concentration has important implications for individuals' economic outcomes. We don't know how the market will evolve. At one extreme, competition and distributed innovation could lead to AI becoming a cheap and ubiquitous commodity. In this scenario, access is democratized and gains are widely shared. Start-ups and smaller businesses have access to state-of-the-art AI resources and can continue their role as a key source of innovative ideas, goods, and services as well as an important engine of job creation for the US economy.

But according to a 2025 paper by Anton Korineck and Jai Vipra, an important factor is that AI has some characteristics that have, in the past, reinforced concentration of market power¹². Like other high-tech innovations, because AI depends on access to data, model improvements, and computing power, it benefits from economies of scale and scope.

Greater data, model improvements, and computing power yield vastly greater intelligence and capabilities. The high return from this advantage helps explain the huge investments and concentration of AI investment in giant firms, which is why they are referred to as 'hyperscalers'.

AI has another feature that seems to be driving ever-greater concentration of market power—the fact that AI itself is a powerful tool to train and accelerate development of new AI models. That is, AI improves its own research and development. While computing technology has always tended to support the market dominance of industry leaders, the extent of the advantage possessed by AI industry leaders may prove to be unprecedented.

As a result of these forces, it is possible that a small number of AI firms may dominate the market and investment returns may accrue primarily to owners of AI¹³. In this potential future, wealth generation—and, to some extent, income generation for those workers who can benefit from access to AI resources—could increase. But less access to ever-improving AI resources for most other firms and their employees would mean slower productivity growth for them, and a steadily widening gap between them and the firms and their employees with more access. As I said at the outset, I am not predicting this particular outcome, just exploring it as a possible scenario.

How AI could alleviate inequality

Let me turn to scenarios in which AI could reduce inequality.

AI as a productivity tool

One scenario is broader access to capability building, resulting in broader productivity gains. Just as the printing press democratized knowledge and the internet democratized information, AI may democratize capability itself. By giving millions of people access to tutoring, coaching, writing assistance, programming support, and problem-solving guidance, AI could enable individuals to develop skills that were once reserved for those with exceptional education, wealth, or mentorship.

In this future, AI could be a rising tide that lifts all boats rather than widening the inequality we've seen in recent decades, sinking other ships. A growing body of research suggests that AI can augment worker productivity without necessarily replacing workers, with especially large gains for less-experienced workers, allowing individuals to perform a broader array of tasks and increasing overall economic output¹⁴.

In one experiment, college-educated professionals completed a range of assignments, writing short reports or analyses, and then carried out a second round of assignments with the help of AI. AI reduced the average time to complete assignments by 40 percent and improved the quality of the results by 18 percent—a combination that constitutes a significant productivity gain¹⁵. The biggest improvements were among those who did the worst on the unaided assignment, narrowing the productivity gap¹⁶.

Acquiring skills or building experience and expertise in any occupation is hard. It demands effort and discipline, often requires years of effort, and depends on education, the capability to acquire and retain knowledge, and the judgment to use knowledge and other resources wisely. AI has the potential to expand expertise, shortening the time and reducing the work needed to build skills, or directly providing those skills themselves, raising productivity the most for lower-skilled, less-educated workers.

AI could also lower the barriers to entrepreneurship for those with a good idea who may lack certain skills, such as finance or accounting, to develop or implement that idea. The democratization of coding capabilities enabled by AI models is also an example of how AI could level the playing field for a critical input for many businesses. While it is too early to know how helpful AI might be to entrepreneurship, the rapid adoption of AI among small businesses is a strong signal of the potential value¹⁷.

New jobs

AI could also create new jobs, ones we haven't even imagined yet, as we've seen with other general-purpose technologies. Consider that by some estimates, there are 12 million full-time social media influencers earning a living today in the United States, something that was unimaginable a decade ago¹⁸. Research suggests that major technological advances lead to this type of job creation, and the bigger the advance, the greater the impact on the labour market¹⁹.

For those who tend to see job dislocation from technological advances as a zero-sum game, it is worth remembering what economists call the 'lump of labour fallacy'. The labour market is not zero sum. For example, spreadsheet software such as Excel replaced the lower-skilled aspects of basic accounting jobs, which raised the productivity of accountants. That is, instead of substituting for a person, technology augmented that person's ability to do their job.

Instead of eliminating a field, technological advances redefined what was possible. AI may be able to do that across the economy. If so, AI could help lower inequality by creating new, more-productive, and higher-paying jobs.

What determines the outcome

I have laid out what are, roughly speaking, the worst- and best-case scenarios of how AI may affect inequality. And

there could be many scenarios in between. It is impossible to predict now which of these versions of the future is more likely, but I will talk through a few examples of factors that could shape that outcome.

Education, job training, and workforce development

The first is education, job training, and workforce development. In the same way that computer skills are essential to many jobs today, proficiency with AI may well be a necessary skill for the jobs of the future.

Workers will need to prompt AI, integrate AI into workflows, oversee coding agents, and manage multiple AI agents. They will need to exercise judgment about AI inputs and outputs, verify AI results, and expand the frontier of human knowledge²⁰. AI's amazing facility in writing computer code is likely to replace individuals whose only job was to write basic code; however, it will empower not only advanced coding experts who will oversee coding agents, but also empower many more people without formal coding training to use coding to turn their ideas into functional programs.

If this is hard to visualize, consider how calculators, word processors, and presentation software have democratized the workplace. Not long ago, these functions were carried out by specialists but are now minimum qualifications for most office jobs.

Education will be critical in how well workers and the US economy adapt to the AI revolution. Economists Claudia Golden and Lawrence Katz explored this history in *The Race Between Education and Technology*, recounting how both the supply and demand for education responded throughout the 20th century to the technological needs of the economy²¹.

First, the dramatic growth in high school completion helped move the country away from agriculture as the dominant industry. In the second half of the century, the proliferation of higher education responded to the

need for specialized expertise in many occupations. In this next wave, a crucial question about how AI may affect inequality will be if high-quality and affordable education and training are widely available not only at the outset, but also throughout one's working life.

Yet a focus on AI skills alone is not likely to be the right approach. My instinct is that curiosity, flexibility, and, importantly, common sense and human judgment are likely to be critical skills in this new economy. As Ethan Mollick argued in his book *Co-Intelligence*, success in the AI era depends less on mastering the technology itself than on developing the human capacities to ask insightful questions²². This includes being able to distinguish sound reasoning from plausible nonsense, make ethical judgments, and integrate knowledge across disciplines.

So, it is not just a question of learning AI skills. To be successful in the future, both young people and those already in their careers are going to need to learn the skills necessary for an economy in which change happens at an increasingly fast pace. Investment in human-centred skills, relationships, and the liberal arts is likely to be as important as technical skills.

Competition and market structure

The market structure of AI firms will also matter to outcomes²³. As explained by Korinek and Vipra, competition is an essential force promoting income and wealth equality. Competition lowers costs, spreads access to technological advances and makes it more likely that the benefits of those advances are shared widely among consumers and workers.

If firms leading the AI revolution achieve dominant positions of market power, the benefits might be concentrated among a fortunate few. Competition alone would not lead to less inequality, as labour and capital markets allocate gains across firms and workers, but competition is an important input into broadly shared gains.

Conclusion

In conclusion, scenarios for AI adoption vary widely regarding how AI might affect inequality. How the market evolves will matter a great deal. But so too will public policy. AI, like past major technological advances, will shape the labour market and the broader economy in myriad ways. It is unclear whether AI will reduce or increase income and wealth inequality, but society can begin making choices now that can affect that outcome.

Decisions on AI policy, education, worker training and workforce development, competition, tax policy, and other areas will help determine this outcome. We have heard many bold pronouncements about what AI will be able to do in the near and distant future. Some will likely come to pass and others won't. But future inequality will depend not only on what AI can do, but on what we choose to do with AI. ■

Michael S Barr is a Member of the Board of Governors of the Federal Reserve System

Endnotes

1. See Michael S Barr, [“What Will Artificial Intelligence Mean for the Labor Market and the Economy?”](#) speech delivered at the New York Association for Business Economics, New York, NY, February 17, 2026.
2. See Gene Sperling, [“An Economic Dignity Compact for the AI Age,”](#) Democracy, January 21, 2026.
3. See [“Income in the United States: 2024,”](#) United States Census Bureau, last modified August 25, 2025.
4. See [“Gini Index,”](#) World Bank Group.
5. See [“DFA: Distributional Financial Accounts,”](#) Board of Governors of the Federal Reserve System, last modified June 18, 2026.
6. See Miles Corak, [“Income Inequality, Equality of Opportunity, and Intergenerational Mobility \(PDF\),”](#) Journal of Economic Perspectives 27 (2013): 79–102.
7. See Claudia Goldin and Lawrence F Katz, *The Race between Education and Technology* (Harvard University Press, 2008).
8. See Board of Governors of the Federal Reserve System, [Economic Well-Being of U.S. Households in 2025](#) (Board of Governors, May 2026).
9. See Maxim Massenkoff, Eva Lyubich, Peter McCrory, Ruth Appel, and Ryan Heller, [“Anthropic Economic Index Report: Learning Curves,”](#) Anthropic, March 24, 2026; and Edward W Felten, Manav Raj, and Robert Seamans, [“Occupational Heterogeneity in Exposure to Generative AI,”](#) April 10, 2023.
10. See Salomé Baslandze, Zachary Edwards, John Graham, Ty McClure, Brent H. Meyer, Michael Sparks, Sonya R. Waddell, and Daniel Weitz, [“Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives,”](#) NBER Working Paper No. 34984 (National Bureau of Economic Research, March 2026); and Jessica Liu and Douglas Webber, [“AI Adoption and Firms’ Job-Posting Behavior,”](#) FEDS Notes (Board of Governors of the Federal Reserve System, March 27, 2026).
11. See Maxim Massenkoff and Peter McCrory, [“Labor Market Impacts of AI: A New Measure and Early Evidence,”](#) Anthropic, March 5, 2026.

12. See Anton Korinek and Jai Vipra, [“Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence,”](#) *Economic Policy* 40, no. 121 (2025): 225–56.
13. See Tania Babina, Anastassia Fedyk, Alex He, and James Hodson, [“Artificial Intelligence, Firm Growth, and Product Innovation,”](#) *Journal of Financial Economics* 151 (2024): 103745.
14. See Ajay K Agrawal, John McHale, and Alexander Oettl, [“Enhancing Worker Productivity without Automating Tasks: A Different Approach to AI and the Task-Based Model,”](#) NBER Working Paper No. 34781 (National Bureau of Economic Research, January 2026); and Shakked Noy and Whitney Zhang, [“Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence \(PDF\),”](#) *Science* 381, no. 6654 (2023): 187–92.
15. See Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, [“Generative AI at Work,”](#) *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942; and Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, [“The Impact of AI on Developer Productivity: Evidence from GitHub Copilot”](#) (2023).
16. See Noy and Zhang, [“Experimental Evidence on Productivity Effects.”](#)
17. See Jeffrey S Allen, [“Monitoring AI Adoption in the U.S. Economy,”](#) FEDS Notes (Board of Governors of the Federal Reserve System, April 3, 2026).
18. See Mark Schaefer, [“How Big Is the Creator Economy? Three Times Larger Than We Thought!”](#) LinkedIn, November 14, 2023.
19. See Daron Acemoglu and Pascual Restrepo, [“The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,”](#) *American Economic Review* 108, no. 6 (2018): 1488–542; and David H Autor, [“Why Are There Still So Many Jobs? The History and Future of Workplace Automation,”](#) *Journal of Economic Perspectives* 29, no. 3: 3–30.
20. See Christian Catalini, Xiang Hui, and Jane Wu, [“Some Simple Economics of AGI”](#) (2026).
21. See Goldin and Katz, [Race between Education and Technology.](#)
22. See Ethan Mollick, [Co-Intelligence: Living and Working with AI](#) (Portfolio, 2024).
23. See Korinek and Vipra, [“Concentrating Intelligence.”](#)

The opinions expressed here are my own and not necessarily those of the Board of Governors of the Federal Reserve System. This article is based on a [speech](#) delivered at 'Next-Gen Financial Inclusion', the third annual Financial Inclusion Conference hosted by the Federal Reserve Board, July 14, 2026.

How the rise of AI matters for fiscal policy

Predictions of the effects of AI on the economy and society tend to focus on productivity growth, job losses, and shifts in the distribution of income. Karen Dynan, Douglas Elmendorf and Louise Sheiner examine various potential scenarios for the macroeconomic effects of AI on US fiscal policy

The rise of artificial intelligence (AI) could have profound effects on the economy and society. Predictions of the economic effects have been offered recently by Acemoglu *et al* (2026), Amodei (2026), Imas (2026), Jones (2026), Karger *et al* (2026), and Trammell and Patel (2025), among many others. Specific predictions vary widely, but three potential effects of AI loom largest for most analysts and commentators: faster productivity growth, job losses, and shifts in the distribution of income.

In our recent paper (Dynan *et al* 2026), we examine four illustrative scenarios for the macroeconomic effects of AI and their implications for US fiscal policy. In the first scenario, AI is a rising tide that lifts all boats proportionally. In the second scenario, AI is again a rising tide, but income gains go only to people in the top quintile of the distribution. In the third scenario, we add significant job losses combined with offsetting increases in labour input elsewhere. And in the fourth scenario, AI causes even faster economic growth, major permanent job displacement, and income gains going exclusively to capital rather than labour.

For each scenario, we form a rough estimate of the effect on the US federal budget in the absence of policy changes, building on the latest projections from the Congressional Budget Office (CBO 2026b) covering the next three decades and on CBO's (2025) analysis of the sensitivity of its projections to changes in economic conditions. We then consider what changes in tax and spending policies might be useful if that scenario were to unfold.

Our analysis points to five conclusions. Taken together, the conclusions show that the fiscal consequences of AI will depend not only on how much it raises national income, but also on who receives that income and how policymakers respond.

Effect of faster productivity growth

First, given current laws and the methodology used in budget projections, faster productivity growth stemming

from AI would improve the budget outlook of the US government – but much of that improvement depends on projected spending patterns that are not consistent with policymakers' past behaviour.

CBO projects that, under current law, US government debt held by the public will rise from 101% of GDP in 2026 to 175% of GDP in 2056. That projection incorporates an AI-induced increase in the annual growth rate of total factor productivity in the nonfarm business sector of 0.1 percentage points over the coming decade; the agency has not publicly quantified its assumptions about AI for later years.

Because the impacts of AI are potentially disruptive but deeply uncertain, policymakers should consider building insurance against the risks

We estimated that, if AI boosts annual productivity growth by a further 0.5 percentage points, the ratio of debt to GDP would be 39 percentage points lower than otherwise at the end of three decades, as shown by the thin solid line in Figure 1. Even so, debt would rise considerably relative to GDP.

The difference in debt stems mostly from a reduction in noninterest spending relative to GDP, which occurs for two reasons. One reason is that payments in benefits programmes are projected to grow more slowly than GDP. In Social Security, this is because the law indexes the benefits of retired recipients to price rather than wage growth; in other benefit programmes, it largely reflects CBO's methodology for projections, which we adopt in our analysis as well. We also assume, like CBO, that faster economic growth has no effect on discretionary spending (the spending that is set through annual appropriations) over the first ten years of the projections.


However, discretionary spending and spending for many benefits programmes have shown no trend relative to GDP over the past three decades or more. If policymakers responded to faster productivity growth by maintaining historical relationships between such spending and GDP, the budget outlook would not improve to the degree that is commonly expected.

Effect of higher income inequality

Second, a shift of income toward the top of the distribution because of AI would further improve the budget outlook if no policies were changed, but it could also spur demands for expanding the tax-and-transfer system to limit the rise in inequality.

An increase in inequality raises federal tax revenue because the tax system is progressive. CBO (2026a) reported that the average federal tax rate was about 21% for all households together and about 26% for households in the top

Figure 1. Federal debt under different scenarios for the rise of AI



Note: Federal debt held by the public relative to GDP. Based on CBO's (2026) extended baseline and authors' analysis.

quintile of the distribution. Accordingly, faster productivity growth with income gains confined to the top quintile of the distribution is projected to leave federal debt 49 percentage points of GDP lower than otherwise after three decades, as shown by the dashed line in Figure 1.

However, if that outcome occurred, it would likely strengthen arguments for greater income redistribution. Those arguments might point to the declining marginal utility of income, the importance of relative consumption for people's wellbeing, and risks to democracy when some individuals acquire disproportionate political power. Increasing the progressivity of federal taxes and spending would have further effects on the budget outlook.

Effect of job losses

Third, job losses caused by AI might not have any automatic effect on the government budget beyond the effects of changes in productivity growth and the income distribution just described, but they would surely generate pressure on policymakers to adopt policies that reduce the harms of such losses.

Workers who lose jobs because AI supplants their line of work would generally struggle to find new jobs that match their skills, so some would be unemployed and others would leave the labour force. The foregone labour input might be made up through additional labour input of some other sort, perhaps from longer workweeks or AI-driven gains in human capital.

Under that assumption, GDP and the distribution of income would be the same as in the previous scenario, and federal debt would again follow the dashed line in Figure 1. Thus, displacement could impose large costs on affected workers without independently changing the government budget outlook.

An extensive research literature has shown that workers who lose jobs often suffer persistent earnings losses, develop significant health problems, and are hurt in other ways. For examples, see Jacobson, LaLonde, and Sullivan (1993), Couch and Placzek (2010), and Finkelstein *et al* (2026).

Policymakers could address these problems using at least three approaches: expanding unemployment insurance, establishing wage insurance, and expanding worker training and employment services. The latter two approaches have been part of the Trade Adjustment Assistance programme, and policymakers could modernise that programme and broaden eligibility to include workers who lose jobs for other reasons. Policy responses along these lines would offset some of the fiscal gains generated by faster growth.

Effect of a large shift of income to capital from labour

Fourth, if an AI-generated increase in national income went entirely to capital owners rather than being shared by capital owners and workers, the budget outlook would improve by less than if the gains were shared proportionally, and policymakers would likely face pressure to adopt policies that broadened the distribution of income gains.

The final scenario in our paper incorporated more dramatic potential effects of AI: an increase in annual productivity growth of 1.0 percentage points (compared with 0.5 percentage points in our other scenarios), major permanent job displacement, and income gains going only to capital. The foregone labour input would be made up, we assume, through additional capital investment, so there would be no effect on GDP or the budget from the loss of labour input itself. But the tax rate on capital income is significantly lower than the tax rate on income generally – by our rough estimate, more than 40% lower – which means that federal debt would be reduced by less.

In our calculations, the beneficial budgetary effects of higher income are just offset by the lower tax rate on that income, and debt would be reduced by 49% of GDP after three decades, as shown by the long-dashed line in Figure 1.

The major job displacement in this scenario would reinforce interest in unemployment insurance, wage insurance, and worker training and employment services that we just discussed. In addition, policymakers could explore ways to subsidise employment in the private sector or provide jobs in the public sector, and they could consider creating a Universal Basic Income programme. Once again, such policy responses would have further budgetary effects.

The shift in income shares from labour to capital would presumably also build pressure to increase the public return from the additional national income. Policymakers could increase the tax rate on capital income – through the existing structure of taxation, by taxing wealth (which might require a constitutional change), or by taxing consumption (presumably with an exclusion for moderate amounts of consumption).

Alternatively, policymakers could establish some public ownership of businesses, which might be similar in simple arithmetic to raising the capital tax rate but could differ in political robustness, sense of shared interests, and opportunities for avoiding payment. That approach could be implemented by creating a sovereign wealth fund or placing equity stakes in individual accounts, either of which would raise a host of complex and consequential policy choices.

Conclusion

Because the impacts of AI are potentially disruptive but deeply uncertain, policymakers should consider building insurance against the risks. For example, they could institute wage insurance and enhanced worker training on

a limited scale to prepare for worker displacement or begin modest public purchases of equities to broaden participation in possible capital-income gains.

Starting modestly would position the government to scale up those efforts later if circumstances warranted. How fiscal policy responds to AI may become one of the central questions of economic policy in our time. ■

Karen Dynan Karen Dynan is Professor of the Practice in the Department of Economics at Harvard University, Douglas Elmendorf is a Professor of Public Policy at Harvard University's Kennedy School of Government, and Louise Sheiner is the Robert S Kerr Senior Fellow in Economic Studies and policy director for the Hutchins Center on Fiscal and Monetary Policy at the Brookings Institution

References

- Acemoglu, D, D Autor and S Johnson (2026), "Building pro-worker AI", The Hamilton Project, February.
- Amodei, D (2026), *The Adolescence of Technology: Confronting and Overcoming the Risks of Powerful AI*. January.
- Congressional Budget Office (2025), *The Long-Term Budget Outlook Under Alternative Scenarios for the Economy and the Budget*. Congressional Budget Office.
- Congressional Budget Office (2026a), Supplement to "The Distribution of Household Income, 2022": Supplemental Data. Congressional Budget Office.
- Congressional Budget Office (2026b), *The Long-Term Budget Outlook Data: 2026 to 2056*. Congressional Budget Office.
- Couch, KA and DW Placzek (2010), "Earnings losses of displaced workers revisited", *American Economic Review* 100(1): 572–589.
- Dynan, K, D Elmendorf and L Sheiner (2026), "How might fiscal policy respond to the rise of artificial intelligence?", NBER Working Paper 35437.
- Finkelstein, A, MJ Notowidigdo and SX Shi (2026), "Trading goods for lives: NAFTA's mortality impacts and implications", NBER Working Paper 34855.
- Imas, A (2026), "What will be scarce?", *Ghosts of Electricity* (Substack blog), 14 April.
- Jacobson, LS, RJ LaLonde and DG Sullivan (1993), "Earnings losses of displaced workers", *American Economic Review* 83(4): 685–709.
- Jones, CI (2026), "A.I. and our economic future", NBER Working Paper 34779.
- Karger, E, et al (2026), "Forecasting the economic effects of AI", NBER Working Paper 35046.
- Trammell, P and D Patel (2025), "Capital in the 22nd century", *Substack blog*, 29 December.

This article was originally published on [VoxEU.org](https://voxeu.org).

A new age of capital: growth, sovereignty and AI

AI is the project that could drive European capital markets forward to channel savings where they are needed, says Christine Lagarde. It needs capital on a scale that only markets can provide, and the stakes for Europe could hardly be higher

Let me begin with a story that starts in Vienna. In Austria, the years before 1873 have a name: the “*Gründerzeit*”, or ‘founders’ era’. The old city walls had come down and the Ringstrasse was rising where they had stood. Railways were pushing out across the empire. New banks and joint stock companies were being founded by the hundred.

But money did not stop at borders. On a scale never before seen, savers in Vienna, Berlin, Paris and London could lend abroad. Their money went into the transformative technologies of the day. French savers financed the Suez Canal. Larger sums still went from British and German savers into the American railways¹.

The historian Eric Hobsbawm called these decades the age of capital². In 1873 that age came to an end. Eight days after the Emperor opened the World Exhibition in the Prater, the Vienna stock exchange collapsed, and four months later Wall Street followed. The crash hit both continents. But what the money had built stayed in the United States, and so did the growth. By the end of the century, the US economy was the largest in the world, while Europe went through two decades of slow growth.

Today, European savers are again financing the transformative technology of the age: artificial intelligence. Again, most of the money is flowing to the United States. Euro area households hold about €440 billion in US technology firms³. The question I want to ask is whether, this time, Europe will get its share of the benefits. I believe it can. But only if we are clear about why we need this technology, why we cannot simply buy it, and why part of it must be built here.

Why Europe must adopt fast

Let me start with some good news. Europe is often criticised when it comes to technology: good at running and regulating an advanced industrial society, but weaker at putting new digital tools to work. With AI, the picture is already different. Things are moving.

Firms are investing significantly. Euro area firms will devote around 10% of total investment to AI in 2026, and ECB staff estimate that AI-related borrowing accounted for about a quarter of the growth in credit to firms in the first quarter of this year. Workers are adopting too. The share of euro area workers using AI on the job has doubled in two years and now exceeds 50%⁴, a level it took the internet about a decade to reach⁵.

Yet progress is not as fast as it could be, as the gap with the United States shows. Over the last two years, US digital investment has grown twice as fast as in the euro area. US workers spend two to three times as much of the working week using AI as workers in the largest euro area economies⁶.

Europe needs a project of its own to drive its capital markets forward. I believe AI is that project. It needs capital on a scale that only markets can provide, and the stakes for Europe could hardly be higher

The reason this matters comes down to growth. An ageing society is straining our social model, and new strategic investment needs are adding to the bill: our analysis finds that existing budgets and EU instruments leave more than €100 billion a year of the public share of those investment needs uncovered⁷. With a workforce that will shrink by more than a million people a year over the next 25 years, there is no way to pay for all this without faster productivity growth⁸.

AI is therefore arriving at the right moment. After years of stalling productivity, it is the best chance we have to make the arithmetic work. ECB estimates suggest that, if adopted quickly, AI could lift the level of productivity by up to 4% over a decade, which would be transformative for public finances⁹. To put that in context, the Draghi report estimated that a 2% rise in productivity over a similar horizon would cover up to one-third of the fiscal cost of Europe's strategic investments¹⁰.

But the case for speed may be even stronger than these estimates suggest, because Europe will pay for this boom whether or not it shares in the growth. AI investment is now growing so fast, above all in the United States, that the tech companies' own cash flows no longer cover it. The big US hyperscalers issued more than USD 100 billion in bonds last year, most of it long-dated¹¹, and they now account for close to a tenth of new euro bond issuance by non-financial companies¹².

Measured against the economy, the buildout is not unprecedented, at about half the size of the IT surge of the late 1990s¹³. But it is arriving at a time when government deficits and debt are high, and market participants already attribute part of the rise in US long-term real yields to AI-related borrowing¹⁴.

Since euro area long-term rates move with US yields, Europe will bear part of the price of this boom in its own borrowing costs. The question is whether it will also get the growth that goes with it.

Why importing is not enough

Adopting, for Europe, means importing. Here, there is an important difference compared with the world of 1873. The railways were a European invention, and it was the United States that was catching up. The *Gründerzeit* earned its name because Europe turned the inventions of that age into companies.

Today, the companies are being built elsewhere. Last year the United States produced 59 notable AI models and China produced 35. France and the United Kingdom produced one each¹⁵. Along the rest of the AI value chain, the picture is the same: the United States hosts three-quarters of the world's AI computing capacity; Europe hosts 5%¹⁶.

There is nothing inherently bad about importing rather than producing new technologies. Historically, the gains from great technologies have mostly gone to those who adopted them, and Europe benefits from US investment in AI. But there are three reasons why AI is special.

The first concerns data. Much of what a firm knows now lies in its data. That knowledge is what sets it apart from its competitors, and economists have started to value it as an asset in its own right¹⁷. Adopting AI means running that data through a system that belongs to someone else and sits under someone else's law. Cloud computing already did this, but a cloud only stores data, whereas a model reads it. That puts the supplier of a model in a position to learn what an industry knows and to compete with it.

European firms understand this, and without domestic alternatives many will hold back. Eurostat asked firms that had considered AI and decided against it why they had done so. Almost half cited data protection concerns¹⁸.

The second reason is access. Past technologies, once adopted, stayed with those who used them. The railway tracks that savers financed in the 1870s could not be taken up again from London or Berlin. Where Europe has remained

reliant on others for inputs or services, the damage a cut-off could cause has been painful but bounded. Halving imports of critical inputs from China would cost us 2-3% of manufacturing value added, for example¹⁹.

With AI, the loss would be of a different kind. Within a few years it will be screening goods at the border, deciding which tax returns are audited, dispatching trains, watching patients on wards and clearing payments at banks. A withdrawal of access, or a change in its terms, would then reach every sector at once. That is leverage of a kind no trade partner has ever held over Europe, and it could be used in any negotiation, on tariffs or on digital taxes, for example.

The third reason concerns the frontier. If we move to a world of many models, including open ones, the threat of a cut-off will diminish. For most of the economy, a model a year behind the best will do the job. But even then the frontier will matter in the most innovative sectors.

Where competition is cut-throat, a small edge in capability can produce outsized gains. Pharma, finance and defence, for instance, all work this way. And this can produce winner-takes-all dynamics: the firm with the best model earns the most from it and can spend the most on the next one, which can entrench its lead for good.

Faced with these properties of AI, Europe has an awkward choice. Either it holds back on adopting, because it cannot protect its data, and forgoes the growth. Or it adopts AI quickly, becomes highly dependent, and risks losing the freedom to organise its economy according to its own values.

Financing sovereign AI

So what can we do to unite growth with sovereignty? The foundation is building more European computing capacity. Companies should be deciding whether to adopt AI on its merits, without worrying about where their data ends up.

There is much debate about whether the AI buildout in the United States is a bubble. ECB staff have looked at past technological revolutions, and corrections have been part of the pattern. No one knows whether this one will follow a similar route, but if it does, we will feel the effects here in Europe²⁰.

Most European savers' exposure to US tech giants runs through investment funds, which would have to sell into a falling market to meet redemptions. Making sure our financial system can absorb that is a pressing task for supervisors and regulators. Banks, which the ECB is responsible for supervising, are well capitalised.

None of this is a reason to delay building on this side of the Atlantic. Europe already has too little data centre capacity to meet its own demand, and on current trends, that gap is projected to grow more than sixfold within a decade²¹. The EU's new gigafactories are a start, but they will fill only a fraction of it. Europe needs to build at a different pace altogether.

Next, Europe needs models that are 'good enough' for most tasks and that run on European infrastructure, so that the threat of being cut off loses its force. Today such models are mostly released with open weights, and the best of them are currently Chinese, such as DeepSeek and Kimi. But they offer no guarantee: a lab can stop releasing new versions openly whenever it chooses, and the licence terms can change with each one.

Europe has something to build on: Mistral in France has released its latest generation of models with open weights, and a firm in San Sebastián has compressed a Chinese open model into a version that now scores higher on independent benchmarks than any other European model. We need to grow that potential to give Europe a range of open models of its own.

The final step is securing access to the frontier, so that Europe stays competitive in the innovative sectors that will drive the fastest productivity growth²². No country controls the whole AI supply chain today. It runs through China for rare earths, Taiwan for chips, Europe for lithography and the United States for models. What countries can do is protect their access by being indispensable to the frontier. Some call this holding a chokepoint. I prefer to see it as being a link in the chain that no one can afford to remove.

Europe holds one of the most important links of all in ASML. But indispensability does not last on its own; others will try to build ways around it. When China restricted rare earth exports in 2010, Japan cut its dependence sharply within a few years. China is now spending heavily to replicate what ASML does. So Europe must keep innovating, to find the links in the chain where it can become indispensable next.

Each step calls for many policies, and I will not list them all. But all of them come back to capital. Europe needs more of it, and it needs a different kind. More of it, because data centres are expensive. The cost of closing Europe's data centre gap over the next decade could be as high as €600 billion, chips included²³.

In the United States the whole of the capital market has been drawn in to make the buildout possible: bonds, private credit and securitisation alongside traditional lending. Europe is still relying largely on its banks. A different kind of capital, because new models and the links in the chain that keep us indispensable cannot be financed with debt alone. They take equity, and equity willing to carry a firm through years of losses before it earns anything.

The news is better than it was. Mistral's funding round last week was the largest equity raise ever by a European technology company. But that is one round, and Europe needs many more. In each case the money is there: European households save around €1.4 trillion a year²⁴. What is missing are the markets that channel it to our own uses.

We hear a great deal about the savings and investments union. This, fundamentally, is what it is for. Savings will always go abroad in search of returns, as they have done since the age of capital. But Europe must first be able to deploy its savings at home.

Let me conclude by stressing the opportunity in front of us. Three years ago I argued that capital markets do not spring up because policymakers call for them. They emerge to finance transformative projects that are beyond the reach of banks²⁵. That is what happened in the 1860s and 1870s, when US railways needed financing on a scale no bank could provide, and capital markets sprung up to channel savings, many of them European, to where they were needed.

I said then that Europe needed a project of its own to drive its capital markets forward. I believe AI is that project. It needs capital on a scale that only markets can provide, and the stakes for Europe could hardly be higher.

Governments across Europe have understood this. The Commission has put its proposals on the table. What remains is to act on them, so that this time Europe's savings can build Europe's future. ■

Christine Lagarde is the President of the European Central Bank

Endnotes

1. Wilkins, M (1989), *The History of Foreign Investment in the United States to 1914*, Harvard University Press; Bonin, H. (2010), *History of the Suez Canal Company, 1858-2008: Between Controversy and Utility*, Droz.
2. Hobsbawm, E (1975), *The Age of Capital, 1848-1875*, Weidenfeld & Nicolson.
3. Chart 3 in Andersson, M, Breckenfelder, J, Corradin, S, Nikolov, K and Viola, MA (2026), [“The AI boom: rational enthusiasm or the next dot-com bubble?”](#), The ECB Blog, ECB, 17 August.
4. Dias da Silva, A, Lebastard, L and Sondermann, D (2026), [“AI adoption and the productivity promise: what workers report”](#), The ECB Blog, ECB, 26 August.
5. Lane, PR (2026), [“AI and the euro area economy”](#), keynote speech at the ECB-SAFE-RCEA International Conference on the Climate-Macro-Finance Interface (3CMFI), Frankfurt, 23 March.
6. Bick, A, Blandin, A, Deming, DJ, Fuchs-Schündeln, N and Jessen, J (2026), [“Mind the Gap: AI Adoption in Europe and the U.S.”](#), NBER Working Paper Series, No 34995, National Bureau of Economic Research, March.
7. Bouabdallah, O, Dorrucci, E, Nerlich, C, Nickel, C and Vlad, A (2025), [“Time to be strategic: how public money could power Europe’s green, digital and defence transitions”](#), The ECB Blog, ECB, 25 July.
8. European Commission (2026), [Demographic transformation in the EU: understanding change and adapting to its consequences](#), Publications Office of the European Union.
9. Lane, PR (2026), op. cit. The ECB scenarios refer to total factor productivity.
10. Draghi, M (2024), [The Future of European Competitiveness](#), European Commission, September, as cited in Bouabdallah, O et al (2025), op. cit.
11. Bank for International Settlements (2026), [“Financing the AI infrastructure boom: on- and off-balance sheet borrowing”](#), BIS Quarterly Review, March.
12. Duquerroy, A, Furtuna, O, Rahmouni-Rousseau, I and Vaz Cruz, L (2026), [“Big tech, big debt: when US tech giants tap the euro area bond market”](#), The ECB Blog, ECB, 31 August.

13. Aldasoro, I, Doerr, S and Rees, D (2026), *“Financing the AI boom: from cash flows to debt”*, BIS Bulletin, No 120, Bank for International Settlements, 7 January.
14. Duquerroy, A et al (2026), *op. cit.*
15. Stanford Institute for Human-Centered AI (2026), [The 2026 AI Index Report](#).
16. Pilz, KF, Rahman, R, Sanders, J, Emberson, L and Heim, L (2025), [“The US hosts the majority of GPU cluster performance, followed by China”](#), Data Insights, Epoch AI, 5 June.
17. Jones, CI and Tonetti, C (2020), *“Nonrivalry and the Economics of Data”*, American Economic Review, Vol. 110, No 9, September, pp. 2819-2858; Veldkamp, L (2023), *“Valuing Data as an Asset”*, Review of Finance, Vol. 27, No 5, pp. 1545-1562; Baley, I and Veldkamp, L (2025), *The Data Economy: Tools and Applications*, Princeton University Press.
18. Eurostat (2026), *Use of artificial intelligence in enterprises, survey year 2025*. The concern has an economic basis independent of privacy law. Handing proprietary data to a model provider creates two risks that a cloud contract did not. First, a provider that sees the workflows of a whole industry is in a position to acquire the knowledge on which its customers’ margins rest and to launch a rival product. Second, once that knowledge is inside the system it can reach other subscribers, through models updated on customer data or through agents and shared features that carry information across firms.
19. Essers, D, Lebastard, L, Mancini, M, Panon, L and Timini, J (2025), [“Critical input disruptions – mapping out the road to EU resilience”](#), Research Bulletin, No 128, ECB, 24 February.
20. Andersson, M, Breckenfelder, J, Corradin, S, Nikolov, K and Viola, MA (2026), [“The AI boom: rational enthusiasm or the next dot-com bubble?”](#), The ECB Blog, ECB, 17 August.
21. According to a study commissioned by the European Commission’s Directorate-General for Communications Networks, Content and Technology, a capacity gap of 3 gigawatts in 2025 rises to around 20 gigawatts by 2036 under the baseline scenario. See Bergamaschi, R et al (2026), *“Cloud & AI Study: Executive Summary”*, European Commission DG CNECT, July.
22. See [“A Transformative AI Strategy for Europe”](#), 14 September 2026.

23. The Commission's impact assessment report for the Cloud and AI Development Act (SWD(2026) 502) projects a gap of 19 gigawatts between EU data centre demand and installed capacity by 2036. Epoch AI estimates the up-front capital cost of a one-gigawatt AI data centre, including servers, at around USD 38 billion (Michael, A and Cottier, B (2026), ["Servers account for 60% of the total cost of ownership of a one-gigawatt AI data center"](#), Epoch AI, 14 May). 19 gigawatts multiplied by USD 38 billion equals around USD 720 billion, or roughly €600 billion at current exchange rates. This should be treated as an upper bound, since the gap covers data centre capacity of all types, while the cost estimate is for GPU-dense AI facilities.

24. European Commission (2025), [Savings and Investments Union: A Strategy to Foster Citizens' Wealth and Economic Competitiveness in the EU](#), 19 March.

25. Lagarde, C (2023), ["A Kantian shift for the capital markets union"](#), speech at the European Banking Congress, Frankfurt am Main, 17 November.

This article is based on a [speech](#) delivered at "Hofburg im Dialog – Economy, Europe, Resilience" in Vienna, Austria, 14 September 2026.

Supercharging AI into explosive growth

AI

If frontier AI models can be used to develop their next models, this may lead to an 'intelligence explosion' and give a large boost to technological advancement. Davidson *et al* develop a framework for understanding how AI may transform the economy through new feedback loops

Frontier artificial intelligence (AI) labs like Anthropic and OpenAI have explicitly stated that they plan to use their AI models to develop their next models, beginning a process of 'recursive self-improvement'. Over time, this may trigger an 'intelligence explosion' as capable models develop even more capable models. If such abundant intelligence is deployed throughout the economy, it may give a tremendous boost to both technological advancements outside of AI and the production of goods and services in the broader economy.

The concept of an intelligence explosion is not new. In 1966, mathematician Irving J Good wrote that *"an ultra-intelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind."*

Yet despite the significant economic implications of such a process, economists have been slow to take the possibility seriously. Aghion *et al* (2019) develop a simple economic environment to re-examine Good's assertion. Suppose a machine's capability grows at a rate that depends on how much capability it has already accumulated, with a single parameter governing the returns to scale of that process. If returns are decreasing — each further increment of capability is harder to win than the last — growth decays toward zero. If returns are increasing, capability diverges to infinity in finite time: an intelligence explosion.

In practice, that parameter may reflect several competing forces. If ideas are getting harder to find, that pushes towards decreasing returns. However, a sufficiently capable machine may conduct research that improves its own capabilities, pushing the other way, and potentially far enough to tip the balance. Whether an explosion occurs therefore depends on which force wins: the drag of diminishing returns or the push of recursive self-improvement.

There have been significant developments in AI in recent years, which help provide a richer view of the intelligence explosion than the single ordinary differential equation offered by Aghion *et al* (2019). In particular, 'scaling laws' suggest that increases in training compute deliver predictable improvements to AI capabilities.

Extrapolating these laws, AI labs have made significant investments in AI chips, with the total stock of compute having doubled every seven months since 2022 (Epoch AI 2026). This ramping up of investment has occurred against the backdrop of predictable technological progress in the capabilities of these chips over the last 60 years ('Moore's Law'), meaning more computing power can be purchased for less.

The question of an intelligence explosion can no longer be dismissed as science fiction – it has become a question of parameters

Algorithmic progress has also contributed to AI capabilities, implying that the training compute required to reach a given level of model performance has declined by about two-thirds per year since 2012 (Ho *et al* 2024).

Given these trends, economic models of AI progress need to account for the growth in the inputs of both compute and software, as well as the distinct research efforts and investment that determine compute stocks and algorithmic progress.


In a recent paper (Davidson *et al* 2026), we develop a framework for understanding the transformative future potential of AI by demonstrating how this technology may establish new feedback loops in the economy. Imagine that AI is powerful enough to perform 20% of tasks in each sector of the economy that used to be performed by workers — including in research.

Adding a stock of 'AI workers' could then increase output both directly — by performing the work required for the production of goods and services — and indirectly — by supporting research and development, ultimately increasing economic productivity. Further, increases in output can then be reinvested in more compute to run and train these AI systems, leading to greater abundance of AI inference as well as better models.


Additionally, there may be a recursive self-improvement feedback loop whereby AI workers can support AI researchers, as well as researchers working on designing more efficient computing hardware. Thus, AI progress may drive additional AI progress. This shift to an economy where 'AI labour' can accumulate through investment is stylised in Figure 1.

Figure 1. AI labour accumulation

Typical economic environment



Automated economic environment



The stylised environment above starts with an initial world on the left where ‘feedback loops’ only occur between capital and output: more output means we can save more to increase the supply of productive capital, which in turn can be used in the production of output.

However, *diminishing* returns to capital mean that this environment only permits a stable balanced growth path, not explosive growth. In the automated economy on the right, however, feedback loops are established in many new places, facilitated by AI which allows capital to replace labour in some tasks.

This automated economy exhibits explosive growth when there are *increasing* returns, ie. when for every dollar the economy allocates to running AI models, it ultimately returns more than a dollar in output. In this environment, economic activity may involve making products, providing services, performing, say, biomedical research, or, of course, researching to develop more advanced AI models.

The question then becomes: what are the conditions in the automated economic environment that permit increasing returns? In our recent paper we attempt to answer this question in a simple environment that integrates the described feedback loops into a macroeconomic model.

We show that in this environment, the only information needed to determine whether the system is explosive is i) the labour share in the economy today, ii) the degree of diminishing returns in technology sectors such as AI algorithms and computing hardware, and iii) the share of tasks in each sector of the economy that are automated. Given estimates of the returns to research effort from Bloom *et al* (2020) and Ho *et al* (2024), we can calibrate this condition. Several insights emerge.

First, if the task bundles across all sectors in the economy are automated equally, the share of tasks that need to be automated to achieve explosive growth is only 13%.

Second, the historical degree of diminishing returns in the hardware sector is much lower than in other sectors. In particular, automating a task in the hardware research sector has about ten times the impact on achieving explosive growth as does automating a task in the production of final goods.

Third, the growth effects of automation are highly convex. The model indicates that 8% of tasks need to be automated in order to double the balanced growth path of output. This is most of the way to the 13% threshold for fully explosive growth.

Although 13% task automation may seem surprisingly low, this figure rests on several important simplifying assumptions. Following Zeira (1998), we assume Cobb-Douglas aggregation of the outputs of tasks completed by AI and human workers. This rules out the possibility that output is bottlenecked by slower growth in the tasks only humans can perform.

By contrast, in an economy with such bottlenecks, automation would generate 'endogenous dampening': large stocks of AI-equivalent workers would simply drive down the elasticity of output with respect to those workers, causing AI-induced growth to trail off (Jones and Tonetti 2026).

We believe bottlenecks of this kind are an important feature of the economics of AI. However, we also show that sufficiently fast automation of new tasks can offset the drag associated with endogenous dampening, thus leaving explosive growth on the cards.

Additionally, even if growth dynamics are explosive, when the strength of feedback links is close to the explosion threshold, it may take many decades before output diverges to infinity. That said, in one illustrative back-of-the-envelope calibration — which abstracts from the bottlenecks discussed above — full software automation in conjunction with a small fraction of automation throughout the rest of the economy could deliver a singularity within six years.

Sixty years after Good's conjecture, the question of an intelligence explosion can no longer be dismissed as science fiction – it has become a question of parameters. Our analysis suggests that the conditions for explosive growth are less demanding than intuition might suggest: modest levels of task automation, particularly in the research sectors that produce better algorithms and better chips, can tip the economy into a regime of increasing returns.

To be clear, this is a statement about possibility. Bottlenecks from tasks that remain stubbornly human, physical constraints on building compute, and the pace at which new tasks can be automated will all shape whether and how fast these dynamics materialise. But the mechanisms are now concrete enough to measure: the labour share, the returns to research effort, and the automation share are all observable.

Economists have historically treated explosive growth as a curiosum at the edge of growth theory. But this is changing: in a recent statement of which all four of us were initial co-signatories (Digital Economy Lab 2026), more than 200 economists and technology leaders — including more than a dozen Nobel laureates — called on the profession and policymakers to *“act now to understand the economics of transformative AI”* and to build the institutions needed to steer it.

Our results underscore why the urgency is warranted: the preconditions for explosive growth are observable and measurable, and may already be in motion. ■

ABOUT THE AUTHORS

Tom Davidson is a Senior Research Fellow at Forethought, Basil Halperin is a PhD candidate in economics at the Massachusetts Institute of Technology, Thomas Houlden is a PhD candidate in economics at Columbia University, and Anton Korinek is Head of Transformative AI Economic Studies at Anthropic Institute

References

- Aghion, P, BF Jones and CI Jones (2019), “Artificial Intelligence and Economic Growth”, in A Agrawal, J Gans and A Goldfarb (eds), *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press: 237–282.
- Bloom, N, CI Jones, J Van Reenen and M Webb (2020), “Are Ideas Getting Harder to Find?”, *American Economic Review* 110(4): 1104–1144.
- Davidson, T, B Halperin, T Houlden and A Korinek (2026), “[When Does Automating AI Research Produce Explosive Growth? Feedback Loops in Innovation Networks](#)”, NBER Working Paper No. 35155.
- Digital Economy Lab (2026), “[We Must Act Now: A Statement on AI’s Transformation of the Economy](#)”, Stanford University, 13 July.
- Epoch AI (2026), “Global AI Computing Capacity Is Doubling Every 7 Months”, Epoch AI Data Insights, published online at epoch.ai.
- Good, IJ (1966), “Speculations Concerning the First Ultraintelligent Machine”, *Advances in Computers* 6: 31–88.
- Ho, A, T Besiroglu, E Erdil, D Owen, R Rahman, ZC Guo, D Atkinson, N Thompson and J Sevilla (2024), “Algorithmic Progress in Language Models”, *Advances in Neural Information Processing Systems* 37 (NeurIPS 2024).
- Jones, CI and C Tonetti (2026), “Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion”, working paper, Stanford GSB and NBER.
- Zeira, J (1998), “Workers, Machines, and Economic Growth”, *Quarterly Journal of Economics* 113(4): 1091–1117.

Authors’ note: Anton Korinek is currently on leave from the University of Virginia and works at the Anthropic Institute, where he heads Transformative AI Economic Studies. This column is based on research performed while he was at the University of Virginia. This article was originally published on [VoxEU.org](#).

Data centres are laying the next financial mine

Yash Kalash argues that the AI infrastructure boom is creating financial transmission mechanisms reminiscent of past crises—not because AI will fail, but because its financing may not endure long enough to repay the infrastructure built around it

The artificial intelligence boom is rapidly becoming an infrastructure boom. Across the United States, Europe, Asia and the Middle East, extraordinary amounts of capital are flowing into data centres, power generation, transmission networks and specialised real estate. Global real estate strategy firm [JLL](#) estimates nearly 100 GW of new data centres will be added between 2026 and 2030, with the sector experiencing what they call an investment supercycle requiring up to USD 3 trillion by 2030.

Much of this investment may prove necessary and AI could become one of the defining general purpose technologies of this century. But regulators should consider another possibility: what happens if technological change moves faster than the investment cycle financing it?

A data centre is financed on assumptions measured in years or decades, while AI technology, model economics and customer demand can change within months. Early expansion was largely funded by [cash-rich hyperscalers](#), whereas [AI labs relied on external capital](#). That balance is changing, for example infrastructure providers such as [Oracle](#) are taking on [long-term investment](#) obligations against contracted AI demand that possibly exceeds their near-term cash generation.

The risk is that debt and long-lived assets are being underwritten against technologies, customers and economics that may change far faster than the infrastructure itself. The Bank for International Settlements has warned that the scale of anticipated AI investment will increasingly require companies to move from internal cash flows toward debt, with [private credit playing a growing role](#).

As the financing ecosystem expands to include banks, private-credit funds, institutional investors and increasingly sophisticated financing structures, technology risk can become financial risk. This creates a familiar equation

reminiscent of the 2008 global financial crisis: optimistic demand forecasts, rapidly appreciating assets, leverage and confidence that the underlying asset will retain its value.

That said, the comparison with the 2008 global financial crisis should not be overstated. Data centres are not subprime mortgages and hyperscalers are not highly leveraged households. Still, the important lesson from 2008 was not simply that American homeowners borrowed too much, but rather that financial markets transformed a reasonable proposition — that housing would remain valuable — into a leveraged financial ecosystem whose stability depended on that proposition continuing to hold.

The dangerous assumption is not that AI will succeed, but rather that today's architecture for delivering it will last long enough to repay everything being built around it

The relevant question today is therefore not whether AI resembles the housing crisis of the past, but whether AI infrastructure is beginning to acquire similar financial transmission mechanisms. The BIS has already identified an AI correction as a potential financial-stability risk, warning that a sharp reduction in hyperscaler capital expenditure could affect borrowers throughout the AI supply chain and ultimately transmit losses into corporate and private-credit markets.

The risk markets may be underpricing is technological obsolescence and evidence is already emerging. [Deepseek](#) in 2025 reported that their V3 model achieved performance comparable to leading closed models while activating only 37 billion of its 671 billion parameters.

Similarly, [Moonshot's Kimi K3](#) activates 104 billion of its 2.8 trillion parameters and claims approximately 2.5 times better scaling efficiency than Kimi K2. And finally, [Apple now runs a roughly 3-billion-parameter model](#) on-device, reducing reliance on centralised compute. Data centres are being built around assumptions concerning today's chips, models, energy requirements and computing architectures.

Those assumptions could change dramatically as models may become substantially more efficient and smaller, move inference onto the edge and new semiconductors, and cooling technologies alter infrastructure requirements, while the geography of computing could increasingly follow abundant and inexpensive energy.

None of this means terrestrial data centres will disappear. The more important possibility is that specialized AI facilities built today may not retain the economic value assumed by their financing models. This deserves particular attention in emerging markets where data centres often depend on scarce long-term finance and substantial supporting investment in electricity grids.

In Malaysia, for example, [IFC provided early-stage financing for a 300MW Johor campus](#) because commercial lenders would not. If utilisation disappoints, the consequences may therefore extend beyond investors to lenders and electricity systems.

Technology companies and infrastructure investors increasingly promote cheaper land, construction costs, expanding electricity supply and skilled workforces as [reasons for governments to facilitate data-centre investment](#).

Some projects will undoubtedly create genuine economic capabilities and governments may offer tax concessions, subsidised infrastructure or scarce electricity based on projections of future AI demand largely originating with the companies seeking capacity. The danger is that host economies socialise infrastructure costs while technology companies retain the flexibility to shift future workloads elsewhere.

Governments should require credible projections of utilisation, technological depreciation, employment, local value creation and long-term energy requirements before committing public resources. At the same time, financial regulators must start looking beyond conventional bank exposure.

They should map who ultimately owns data-centre risk across banks, private-credit funds, insurers, pension funds and securitised products. Stress tests should examine what happens when hyperscalers reduce capital expenditure, utilisation assumptions disappoint and specialised real-estate valuations decline simultaneously.


Crucially, scenarios should also include a world in which AI demand continues growing while the amount or type of infrastructure required to provide it changes dramatically — AI does not have to fail for an AI infrastructure bubble to burst.

History provides ample precedent. Railways transformed nineteenth-century economies while [repeatedly bankrupting railway investors](#). The internet transformed the world despite destroying enormous amounts of capital [after the dot-com bubble](#). During the Great Financial Crisis housing still remained indispensable even as securities constructed around assumptions about its value destabilised the global financial system.

A transformative technology and a speculative investment cycle can often coexist and thus regulators should begin treating the gap between the speed of technological change and the duration of infrastructure investment as a potential source of macro-financial risk.

The dangerous assumption is not that AI will succeed, but rather that today's architecture for delivering it will last long enough to repay everything being built around it. ■

S Yash Kalash is a Senior Fellow at CIGI and an expert in strategy, public policy, digital technology and financial services

The background of the slide is a photograph of the Bank of England building, featuring its iconic neoclassical architecture with large columns and a pediment. The image is darkened to serve as a backdrop for the text.

Some questions that a BoE review should ask

As the Bank of England nears 30 years of operational independence, a broad review is likely. Sushil Wadhwani proposes five measures to reduce future quantitative-easing losses and, as markets still doubt the Bank will meet its inflation target two years after the Bernanke Review, steps to strengthen credibility

The independence of the Bank of England (BoE) is perhaps the best thing that has happened to British macroeconomic policy in my lifetime. In recent years, though, parliamentary committees have been especially interested in quantitative easing (QE) (eg. House of Lords 2021, House of Commons 2024) and, a few weeks before being appointed as the First Secretary of State, Louise Haigh, argued that the time was right to re-examine the BOE's mandate (Haigh 2026). The reforms that I propose below should help deal with some of this political criticism and thereby entrench independence.

What might we do regarding the losses associated with QE

The Office for Budgetary Responsibility's most recent published estimates suggest that the cumulative net lifetime loss from the BOE's Asset Purchase Facility (APF) might amount to £164 billion by the end of 2036 (OBR 2025, para 6.18).

While the OBR emphasises that this excludes the wider benefits of the QE programme, it nevertheless shows that a government that is committed to its fiscal rules ends up probably having to spend less or tax more than might have been true in the absence of these losses.

Recall that the procedures associated with QE were put together at a time of great stress for all involved and it would be reasonable to expect that these can be enhanced. I therefore make five recommendations below with regard to how we might reduce these APF losses in the future, but it is best to understand the underlying problem first.

Consider the stance of the Monetary Policy Committee (MPC) that it pays no attention whatsoever to the profit or loss associated with the assets transacted when making their monetary policy decisions. This is because the government indemnifies the BOE against future losses. One justification offered by the BOE is that when it buys

a gilt, its expected return on the overall transaction (which typically involves financing the gilt purchase through paying interest on reserves) is approximately zero because markets are efficient and accord with the expectations theory of interest rates¹.

In this regard, Shiller (1979) showed that long-term bond yields move around substantially more than the conventional expectations theory of the yield curve can justify. An implication is that bond returns are somewhat predictable and so, when buying a longer-term gilt, there are some times when you want to be a buyer and others when you wish to be a seller.

The two suggestions made to enhance the credibility of the BoE and the inflation targeting regime should both help reduce the debt interest bill and the cost of reducing inflation. The recommendations with regard to reducing QE losses should hopefully make central banking more boring and reduce the degree of unwanted attention from politicians

It is important to remember that central banks have typically launched QE in difficult situations involving a financial crisis. Think of them as periods when 'animal spirits' are unusually low, as perhaps evidenced by stock prices being 'irrationally' low too. Investors typically flee to the safe haven provided by longer-term bonds and this depresses the term premium.

So, central banks know that if and when the economy normalises, interest rates and the term premium will both rise and bond prices will fall. It is therefore likely that QE will ultimately lose money as it involves 'buying high, selling low'.

The first recommendation is that, given that accountability is important, the MPC should revisit its efficient markets view and publish scenarios contemporaneously showing that significant APF losses might materialise and why the Committee nevertheless believes that the QE operations are still likely to deliver a net benefit to the economy as a whole. The very act of having to be explicit in this regard should help focus minds on the MPC and exert downward pressure on losses.

The BOE also justifies its stance on ignoring the expected losses associated with QE by saying it needs to exclusively focus on achieving the inflation target and not be distracted by fiscal policy considerations. But the MPC can achieve the inflation target in multiple ways, and differences in the level of expected losses can help them choose between these alternative strategies.

For example, the MPC could choose to buy UK equities instead of gilts during these periods of panic in order to help achieve its inflation target (the Bank of Japan also bought stocks in recent years). An advantage of doing so is that buying equities would have made money.

So the second recommendation is that the MPC should discuss and show comparisons in likely costs between alternative ways of achieving the inflation target (eg. buying equities versus gilts, buying gilts at different points along the curve, etc). I might note here that Australia, in implementing the suggestions of a review of its central bank, has agreed that the Monetary Policy Board will include the benefits, costs, and risks associated with available tools when contemplating their use².

A third recommendation is to facilitate greater policy coordination between the BOE and the Treasury. The MPC would, initially privately, share its analysis of its particular form of proposed QE stimulus, showing the expected fiscal costs in comparison to the expected benefits (as per the recommendations above). The Treasury would then be given the option to propose an alternative fiscal policy measure which provides a superior benefit/cost ratio while also helping to deliver the inflation target. Minutes of these discussions would be published with a time lag.

An exception would need to be made for situations where the MPC believed that action on its part was very urgent. Of course, the details of such an arrangement deserve further consideration as while it would be desirable in terms of ensuring that the return for a given fiscal cost is optimised, it must be compatible with the BOE being seen to be independent.

A fourth recommendation is that the MPC should revisit whether the QE instrument has become unnecessarily inflexible because of the Committee's tendency to pre-commit to purchases over a long period. For example, the MPC committed itself on 4 November 2020 to buy £150 billion of gilts over the subsequent year, and yet on 9 November, highly encouraging vaccination results were announced.

If, as per earlier suggestions, the MPC had been told to consider the expected profits/losses associated with QE in conjunction with the overall economic benefit, it would probably have truncated the pre-announced QE programme after the vaccine news.

A fifth recommendation is to reduce the fiscal costs of QE by operating a tiered reserves remuneration structure (as the ECB does; see De Grauwe and Ji 2024). Of course, the details of the change here should pay due consideration to the need to preserve monetary policy transmission.

How Bank of England reform might help strengthen credibility

I have previously discussed on VoxEU that the level of breakeven inflation embedded in UK gilt yields was high in relation to its inflation target and international peers (Wadhwani 2024a). This remains true today³. The 20-year breakeven inflation rate in the UK is currently around 3.32%, which is high in relation to the inflation target of 2%.

We do though need to account for a quirk whereby UK index-linked gilts compensate investors for Retail Prices Index (RPI) inflation for the first four of these twenty years, and Consumer Price Index including occupiers' housing costs (CPIH) inflation from 2030 onwards. Note that the four-year ahead breakeven inflation rate is approximately 3.85%.

This then implies that the expected average inflation compensation for the 16-year period beginning in four years that is embedded in markets is around 3.19%, which is well above the inflation target of 2% and is broadly in line with the level of 3.17% reported in Wadhwani (2024a).

By comparison, the level of breakeven inflation in the US is around 2.40% (for CPI inflation), which is broadly in line with the Federal Reserve's target of 2% for Personal Consumption Expenditures (PCE) inflation. There is therefore a

significant 'inflation risk premium' embedded within UK gilts, while investors in US Treasuries do not appear to need such a risk premium.

One possible reason for the higher inflation risk premium in the UK is that the government sets the inflation target while, in the US, the Federal Reserve chooses its own definition of price stability (ie. goal independence). I am tempted to recommend that the BOE be allowed to set its own target, thereby bringing it in line with the US and the ECB.

But some have argued that the operational independence of the BOE was only feasible because Parliament retained control of the inflation target (Stansbury and Balls 2017). Therefore, an alternative way forward is to enshrine the inflation target into law using primary legislation, as this would still signal that the government is serious about entrenching price stability.

Another possible reason as to why the inflation risk premium is elevated in the UK is that the BOE itself is perceived as less credible than some other central banks. The proposals within the Bernanke Review (Bernanke 2024) were regarded as a way to help enhance BOE credibility, but as I argued at the time (Wadhvani 2024b), the forecasting errors made by the Bank had much more to do with a lack of intellectual diversity.

By way of example, the Bank had failed, at the end of 2020, to notice that its forecasts were incompatible with the ongoing money growth and most reasonable beliefs regarding the behaviour of velocity. It therefore does not surprise me that the level of the inflation risk premium has remained broadly stable in recent years even after the implementation of some of the recommendations of the Bernanke Review, as the markets appreciate that more needs to change⁴.

The fact that inflation has spent much of the last five years above target has not helped. Specifically, inflation moved above target in May 2021 and has been above target in all but three months since then (it was at the 2% target in May-June 2024 and below target in September 2024). Therefore, I recommend that the review looks afresh at recruitment procedures to ensure that we get a more intellectually diverse membership.

Haigh (2026) has suggested that the BOE remit be amended to pay more attention to growth by, perhaps, switching to a dual mandate. Recall that the BOE's statutory objective is to maintain price stability and, subject to that, support the government's economic policy (including growth). Moreover, the remit explicitly allows the MPC to return inflation to target gradually when immediate disinflation would generate excessive output losses.

So, I am yet to be convinced that switching to a dual mandate would make much difference in practice to how the MPC actually sets interest rates. On the other hand, there is a risk that it would be misunderstood by markets as a signal that the government is less committed to the inflation target and might thereby widen the inflation risk premium.

Conclusion

The two suggestions made to enhance the credibility of the BoE and the inflation targeting regime should both help reduce the debt interest bill and the cost of reducing inflation. The recommendations with regard to reducing QE losses should hopefully make central banking more boring and reduce the degree of unwanted attention from politicians. ■

Dr Sushil Wadhwani, CBE, is currently a Research Associate at the London School of Economics

Endnotes

1. See, for example, Ben Broadbent's [response](#).
2. See Mahon and Milas (2026) for further discussion.
3. Using data on 10 September, 2026.
4. On the day before the publication of the Bernanke Review (11 April 2024), the estimated six-year-ahead, fourteen-year breakeven inflation rate was 3.23%, which compares to our estimate of underlying long-term inflation expectations of 3.19% now.

References

- Bernanke, B (2024), *Forecasting for monetary policy making and communication at the Bank of England: a review*, Bank of England.
- De Grauwe, P and Y Ji (2024), [“The new operating procedures of the Bank of England”](#), VoxEU.org, June 24.
- Haigh, L (2026), *“A new fiscal framework to renew Britain”*, *Renewal* 34, May.
- House of Lords (2021), *“Quantitative easing: a dangerous addiction?”*, July 16.
- House of Commons (2024), *Quantitative Tightening, fifth report*, January 31.
- Mahon, C and C Milas (2026) *“Lessons in letter writing for the Bank of England”*, Alphaville, *Financial Times*, 10 August.
- OBR – Office for Budgetary Responsibility (2025), *Economic and fiscal outlook*, November 2025.
- Shiller, R (1979), *“The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure”*, *Journal of Political Economy* 87(6): 1190-1219.
- Stansbury, A and E Balls (2017), [“Twenty years on: Is there still a case for Bank of England independence?”](#), VoxEU.org, 1 May.
- Wadhwani, S (2024a), [“How to cut the UK debt interest bill”](#), VoxEU.org, 22 July.
- Wadhwani, S (2024b) *“The Bernanke Review: A Missed Opportunity?”*, in D Aikman and R Barwell (eds), *The Bernanke Review: Responses from Bank of England Watchers*, King's Business School.

Editors' note: The author is a former member of the Monetary Policy Committee at the Bank of England. This article was originally published on [VoxEU.org](https://voxeu.org).

From vision to delivery

The ECB wants central bank money to serve as the monetary anchor for tokenised markets. Piero Cipollone calls for common standards, central bank money and coordinated regulation to prevent blockchain-based finance from becoming a collection of disconnected platforms

Two years ago, I argued that Europe had a unique opportunity¹. Tokenisation and distributed ledger technology, or DLT, were beginning to reshape financial markets. By representing and transferring assets in the form of programmable data files – tokens – they offered the prospect of making finance more efficient, enabling it to operate around the clock with greater automation and fewer intermediaries.

At the heart of this transformation lies a Copernican revolution. Tokenised assets will be at the centre of a new digital finance ecosystem that allows issuers, investors and intermediaries to think anew and innovate. But this transformation also poses a risk: a proliferation of incompatible platforms could reproduce, or even deepen, the current fragmentation of Europe's capital markets.

I therefore outlined a vision for building an integrated and dynamic European ecosystem for digital assets from the outset, with central bank money at its core. Since then, we have moved from vision to delivery. And I will discuss the next steps in turning that vision into reality.

I will first recall the objectives we set ourselves and why they remain our cardinal points. I will then explain the progress we have made, in particular through our Pontes and Appia projects. Finally, I will outline the conditions for success, namely what it will take to turn technological change into a genuine transformation of finance in Europe.

The vision: an integrated European market for digital assets

The promise of tokenisation is not simply that existing processes will become faster, but that it will enable us to simplify the architecture of finance itself. This is particularly relevant in Europe, where capital markets remain deeply fragmented despite some recent progress towards integration².

The EU now has 31 central securities depositories (CSDs), 14 central counterparties (CCPs) and 323 trading venues³. And while some larger groups cover the value chain – encompassing trading, clearing, settlement and asset servicing – crossborder settlement remains limited. Securities are still predominantly held and settled within local CSDs.

Even in CSDs belonging to the same group, most transactions are settled within individual CSDs in the group, with limited crossborder transactions. In 2023 more than 95% of transactions in both volume and value terms were settled between parties in the same individual CSD⁴.

If we design and build an integrated European market for tokenised assets from the outset, the digital finance transformation will allow us to leapfrog the fragmentation of existing legacy systems

So how can tokenisation help simplify our financial ecosystem and make it more efficient? Today, the lifecycle of a financial asset often involves a succession of infrastructures and ledgers. Issuance, trading, clearing, settlement, custody and asset servicing may all take place in different systems. Each step requires information to be transmitted, verified and reconciled.

Tokenisation could bring several of these functions into a shared digital environment. Transactions could be executed conditionally and atomically, meaning that the cash and asset legs are either settled together or not at all. Smart contracts could automate processes ranging from coupon payments to collateral movements and compliance checks.

Tokenisation is therefore not merely an incremental improvement to one stage of a financial transaction. Instead, it has the potential to reorganise the entire financial value chain. This reorganisation provides an opportunity for Europe. If we design and build an integrated European market for tokenised assets from the outset, the digital finance transformation will allow us to leapfrog the fragmentation of existing legacy systems.

But technology alone will not determine the outcome.

The same technology can support an integrated market or a collection of closed platforms. It can lower barriers to entry or create new gatekeepers. It can foster competition or entrench technological dependencies.

This is no longer a purely hypothetical question. Globally, tokenised finance is moving from experimentation to early adoption, as institutional players increasingly focus on tokenising traditional financial instruments. While still small, the figures point to genuine momentum.

Worldwide, tokenised traditional assets recorded on public blockchains increased roughly fivefold between March 2025 and March 2026⁵. In some segments, activity has already reached meaningful operational scale. In the United States, one private platform processed an average of USD 354 billion in tokenised repo transactions per day in March 2026, four times the average daily volume it recorded a year earlier⁶.

In Europe, adoption is progressing and European institutions are developing tokenised bonds, deposits, collateral and settlement solutions. European CSDs have also announced major initiatives to tokenise securities on a large scale, matching similar initiatives in the United States. And in March, we started accepting marketable assets issued via DLT-based services at European CSDs as eligible Eurosystem collateral⁷.

Compared with the size of global markets, tokenised real-world assets remain tiny, with limited liquidity and secondary market activity. But we are at a critical juncture. The market is sufficiently advanced for the opportunities to be tangible, but still young enough for Europe to shape its architecture and overcome coordination challenges.

In doing so, we need to avoid three risks. The first is fragmentation. Incompatible networks could potentially divide assets, liquidity and participants into separate ecosystems. This would disperse activity across platforms that may not communicate with each other, fragmenting liquidity.

The second risk is the loss of the monetary anchor. Without access to central bank money, tokenised markets could become reliant on private settlement assets which carry credit and liquidity risks and do not have the same ability to scale elastically.

The third risk is external dependence. Europe could become reliant on infrastructures, technologies, governance arrangements or settlement assets controlled outside Europe.

The question is therefore no longer whether tokenised finance is emerging, but whether it will develop as an integrated and competitive market in Europe, or as a collection of disconnected technological islands that fail to scale up and compete globally.

The Eurosystem's objectives are clear. We want to offer settlement of DLT transactions in central bank money, providing a safe asset and a monetary anchor on which tokenised finance can grow. We want to foster an integrated, competitive and innovative European payments and securities ecosystem, which would also support the euro's attractiveness globally. And we want to preserve Europe's resilience and strategic autonomy⁸.

The progress: from experimentation to implementation

Two years ago, I argued that public authorities needed to act early, before fragmentation became entrenched. We have acted, and we have moved at pace.

The EU co-legislators are discussing the legislative proposal to extend and enhance the DLT Pilot Regime. At the same time, the European Commission has launched a consultation to assess whether and how the Markets in CryptoAssets Regulation should be fine-tuned, drawing on the experience gained in its first years of application.

The Eurosystem is also making quick progress. Its first step was to assess the market's needs and test possible technological solutions. In 2024 we brought together 64 market participants in more than 50 trials and experiments. This work showed that central bank money can be used to settle transactions conducted on DLT platforms. It also confirmed that access to central bank money is not a secondary consideration for the market, but one of the conditions for tokenised finance to develop safely and at scale⁹.

We have transformed those findings into a coherent strategy, which we are implementing through two complementary projects: Pontes and Appia. Pontes will turn our commitment to provide central bank money for settling tokenised transactions into an operational service. It will connect market DLT platforms with the Eurosystem's TARGET Services. Final settlement of the cash leg will take place in central bank money. And synchronisation will enable delivery versus payment and other transactions that require all-or-none settlement.

In 2024 we said that tokenised markets needed access to central bank money. This year, we intend to deliver on what we said by going live with Pontes. To support its early adoption, we will offer attractive pricing conditions, charging only one-off onboarding fees for the initial launch¹⁰.

We will then gradually enhance the service by extending operating hours to 22.5 per business day and offering immediate settlement finality in the Eurosystem DLT¹¹. By mid-2028, we plan to offer a 24/7 service as well as greater programmability, state-of-the-art resilience and multi-currency capability.

Pontes matters, because confidence in a settlement asset is not something that technology can manufacture. Central bank money does not carry credit or liquidity risk. What's more, it serves as the common settlement anchor across the financial system.

By making central bank money available for DLT-based transactions and bringing central bank money on-chain, Pontes will provide market participants with a firmer basis for investing, innovating and scaling up. At the same time, it will generate valuable operational experience that can be taken into account when designing the long-term ecosystem.

But providing central bank money is a necessary, not a sufficient condition for expanding and integrating digital finance in Europe. That's why we launched Appia. Appia addresses the broader architecture, standards and governance of a European tokenised financial ecosystem. The roadmap we published in March combines analytical work with practical experiments and close cooperation between the market and public authorities.

The work conducted as part of the Appia roadmap covers asset interoperability and standards, monetary policy implementation and collateral management, the future infrastructure for tokenised central bank money, crossborder transactions, as well as the ecosystem's legal and regulatory framework, and foundation for resilience. It aims to deliver a blueprint for an integrated, European tokenised financial ecosystem in 2028¹².

Put simply, Pontes builds bridges by offering digital finance a safe settlement asset and by making private settlement assets mutually convertible – enabling, for example, tokenised deposits to be transferred between banks or stablecoins to be settled in fiat currency directly via DLT.

Appia paves the way for a single European market for tokenised assets, but we have not predetermined what the architecture for that will actually look like. On the one hand, a single shared European network would avoid fragmentation and would mutualise infrastructure costs.

On the other hand, multiple interconnected networks would support innovation and resilience at the infrastructure level. However, these networks would require a high degree of interoperability to avoid the fragmentation of assets and liquidity. Combinations of these approaches would also be possible.

Appia will therefore assess the alternatives against our objectives for integration, competition, resilience, innovation, governance, strategic autonomy and the ability for the Eurosystem to set and enforce the rules applicable to central bank money.

Pontes and Appia should not be seen as separate projects, but as two parts of one strategy. Appia's findings will feed into the gradual enhancement of Pontes. At the same time, the operational experience gained from Pontes will inform the choices made under Appia.

We are combining immediate delivery with long-term ambition. We are learning by doing, while ensuring that each operational step takes us closer to a dynamic, integrated and competitive European market.

The path to scale: the conditions for success

Interoperability based on common standards and rules

The first condition is interoperability across DLT platforms based on common standards and rules. This will enable investors to access the vast majority of tokenised assets and services by joining only one or a few DLT networks. For this reason, the revised DLT Pilot Regime proposed by the Commission requires the industry to establish technical standards that support interoperability between DLT market infrastructures.

To achieve interoperability, connecting two ledgers is not enough. The assets must mean the same thing on both ledgers, rights must still be enforceable, the rules set by issuers must still apply, and transfers must achieve legal and operational finality.

A genuinely integrated tokenised ecosystem must therefore have five capabilities.

- Systems must be interoperable so that they can exchange and understand trusted instructions and transaction states.
- Assets must be transferable, with valid authorisation, compliance and settlement finality.

- They must also be portable, and their identity, rights, obligations and history must be preserved when they move between platforms.
- They must remain controllable, so that issuers and public authorities can exercise the powers required in regulated markets.
- Lastly, they should be programmable, allowing automation to develop within a safe, legally valid and governable framework.

Achieving these capabilities requires more than technical protocols. It requires coordination across infrastructure, identity, data, asset representation, transaction mechanisms, governance, risk controls and supervision.

Europe's experience with SEPA and TARGET Services teaches us that integration requires common standards, harmonised practices and shared rules.

We should build on mature industry and international standards where they exist, rather than duplicate them, but we must also ensure that they are adopted consistently and fill in the remaining gaps.

Our immediate priorities should include mechanisms for transferring cash and securities tokens between platforms, compatible standards for representing those tokens, and common reference data that allow assets and lifecycle events to be interpreted consistently.

This does not mean imposing a one-size-fits-all on the market. Instead, it means ensuring that different solutions can interact and compete on the same market. This is an essential condition for tokenised finance to achieve

scale and be widely adopted, as was the case for electrification, when the associated efficiency gains relied on the regulatory and standardisation choices made in the early 20th century¹³.

There should be competition in services, quality and price, not through incompatible standards or walled gardens.

A public-private partnership to drive adoption

The second condition is public-private cooperation that will lead to coordinated adoption. Tokenisation is a system-wide innovation. Its benefits will only fully emerge if issuance, trading, settlement, custody, collateral and asset servicing all develop alongside each other.

Issuers need investors and liquid secondary markets. Investors need reliable custody, settlement and asset servicing. Market-makers need effective repo and collateral arrangements. And all participants need legal certainty and access to a trusted settlement asset.

If each actor waits for the other to play their part, tokenisation may remain confined to successful but isolated pilot schemes. The public sector can provide the monetary anchor and establish a supportive European regulatory framework. But the market must provide the assets, the services and the innovation.

Together, we should not just coordinate experiments, but aim to create the conditions for widespread adoption, making it possible to reap complementarities across investors, issuers and intermediaries and enabling the market to scale up.

This is a joint endeavour that involves market participants, public authorities and members of academia. In Appia, we will jointly conduct targeted experiments, develop and test standards, and pursue practical work on asset transfers and collateral mobility. Throughout this process, we will seek continuous feedback from Pontes users.

An integrated legal framework

The third condition is a legal framework matching our technological ambition. Advanced technology cannot compensate for fragmented law.

If an asset can move technically from one platform to another but its ownership status becomes uncertain in the process, it is not genuinely portable. If a transaction is synchronised in code but its finality is not recognised across jurisdictions, it is not genuinely settled.

Technical interoperability without legal compatibility will remain incomplete and fail to overcome fragmentation. While the enhanced DLT Pilot Regime requires standards to ensure technical interoperability and seeks to open the market to new players, we need to provide the legal foundations for tokenised assets to be issued, held and transferred seamlessly across the EU.

Europe therefore needs greater clarity and harmonisation on the legal status of tokenised assets, ownership rights, settlement finality, liability, custody, asset servicing and the enforceability of smart-contract outcomes. Otherwise, we risk building advanced digital rails on fragmented legal foundations.

The three conditions I have described are closely connected. Common standards make interoperability possible. Public-private cooperation turns interoperability into adoption. And legal certainty allows adoption to occur safely and across borders.

Conclusion

Two years ago, I argued that Europe had a unique opportunity to build an integrated capital market for digital assets from the outset. Today, we are delivering on that vision.

Pontes is turning the settlement of tokenised transactions in central bank money into an operational reality. Appia is turning a long-term ambition into a public-private roadmap.

But the next step will be decisive. We need to turn technological connectivity into genuine market integration. That requires common standards, interoperable assets, open access, effective governance and legal certainty. To create these conditions, public authorities and market participants need to work together.

If we wait too long to take action, fragmentation may become entrenched. If we act now, we can do more than just modernise settlement. We can use technological change to help build something Europe has sought for decades: a deeper, more innovative and more competitive capital market.

The vision is clear. The foundations are being built. And on those foundations we can build the future of European finance. Together. ■

Piero Cipollone is a Member of the Executive Board of the European Central Bank

Endnotes

1. Cipollone, P (2024), [“Towards a digital capital markets union”](#), keynote speech at the Bundesbank Symposium on the Future of Payments, Frankfurt am Main, 7 October.
2. ECB (2026), [Financial integration and structure in the euro area](#), 7 May, and Lagarde, C (2024), [“Follow the money: channelling savings into investment and innovation in Europe”](#), speech at the 34th European Banking Congress: “Out of the Comfort Zone: Europe and the New World Order”, Frankfurt am Main, 22 November.
3. European Securities and Markets Authority (2025), [ESMA statistics on securities and markets](#), 12 December.
4. Hanssens, B, Sandín de Vega, D and Sowa, HF (2025), [“Advancing the capital markets union in Europe: a roadmap for harmonising securities post-trading”](#), Economic Bulletin, Issue 4, ECB.
5. From €4.7 billion at the end of the first quarter of 2025 to €23.3 billion at the end of the first quarter of 2026.
6. The activity of Broadridge’s Distributed Ledger Repo platform corresponds to roughly 7% of daily US repo activity. However, this comparison needs to be interpreted with caution to the extent that activity on the Broadridge DLR platform includes intraday transactions, implying a higher turnover velocity than in conventional repo statistics.
7. ECB (2026), [“ECB paves way for acceptance of DLT-based assets as eligible Eurosystem collateral”](#), press release, 27 January.
8. ECB (2026), [Appia – paving the way for a future-ready, integrated financial ecosystem leveraging tokenisation and DLT](#), 11 March.
9. ECB (2025), [Bridging innovation and stability: the Eurosystem’s exploratory work on new technologies for wholesale central bank money settlement](#), June.
10. ECB (2026), [Pontes Pricing Guide](#), 19 August.
11. Initially, settlement finality will be achieved once the corresponding movement of funds is reflected in T2.
12. ECB (2026), [Appia – paving the way for a future-ready, integrated financial ecosystem leveraging tokenisation and DLT](#), 11 March.

13. Cipollone, P (2026), *"Sparking the transformation of finance: tokenisation and the role of central banks"*, keynote address at the 24th Annual Symposium on "Building the Financial System of the 21st Century: an Agenda for Europe and the United States", hosted by the Harvard Law School and the Program on International Financial Systems, Washington DC, 15 April.

I would like to thank Jean-Francois Jamet for his help in preparing this [speech](#), as well as Thomas Vlassopoulos and Alessandro Speciale for their comments. This article is based on this speech, delivered at the Deutsche Bundesbank's Symposium on 'Future of payments: trends and innovations in Germany and Europe', Frankfurt am Main, 26 August 2026.

Achieving a European Banking Union without increasing risks

European banking remains fragmented, leaving billions of euros inefficiently allocated. Gyöngyi Lóránth, Anatoli Segura and Jing Zeng argue that a common deposit insurance system can encourage more risk-taking, but this should be accompanied by sufficiently strict supervision

More than a decade after its creation, the banking union remains incomplete. Supervision has been centralised, and a common resolution framework has been created, but deposit insurance remains national. This incomplete allocation of responsibilities has been linked to the continued fragmentation of European banking, with few crossborder mergers and limited crossborder competition (Schnabel and Veron 2018, Beck *et al* 2022).

As Andrea Enria, then Chair of the ECB's Supervisory Board, put it in 2023, *"an incomplete banking union is the reason why crossborder banking groups are ring-fenced along national lines and crossborder integration does not happen"* (Enria 2023). According to ECB estimates, national restrictions leave €225 billion of capital and €250 billion of liquidity trapped in local subsidiaries. These resources cannot necessarily be deployed where they are most needed, making crossborder groups less efficient and crossborder expansion less attractive (Financial Times 2026).

In June 2026, the European Commission was reported to be considering reforms that would give banking groups greater freedom to allocate resources across borders, while requiring parent banks to support subsidiaries when needed. The European Commission Executive Vice-President Teresa Ribera urged member states to support crossborder bank mergers and pointed to the absence of a common deposit-guarantee system as one of the obstacles to completing the banking union (Reuters 2026). A common European deposit insurance system would reduce the incentive to protect resources along national lines and allow crossborder banking groups to operate more like genuinely integrated firms.

Missing in this policy debate is a key moving part: how do banks change their risk-taking behaviour once resources can move freely within the group, and particularly in a crisis? Removing ring-fencing can affect bank risk-taking because it changes not only the allocation of resources after trouble occurs but also the consequences of getting into trouble in the first place. A subsidiary that expects support from and extends support to the rest of its group faces a different set of incentives from one that expects each unit to recapitalise itself.

In our research, we show that this link between *risk-sharing* and *risk-taking* fundamentally changes the welfare consequences of completing the banking union (Lóránth *et al* 2026).

Why national authorities ring-fence

Consider a crossborder banking group with a healthy subsidiary in one country and an impaired subsidiary in another. From the group's perspective, the natural response is to transfer resources from the healthy unit to the impaired one. Such support may prevent costly external recapitalisation and preserve the value of the group.

Greater freedom to move capital and liquidity across borders should be accompanied by sufficiently strict supervision to address the risk-taking incentives created by expectations of future within-group support

The national authority responsible for the healthy subsidiary sees a different problem. If resources leave the country and the subsidiary later fails, fewer resources remain to protect local depositors. A transfer may therefore reduce the expected losses of the banking group as a whole while increasing the losses borne by one national deposit insurance fund. This is the fundamental source of ring-fencing: When responsibility for losses is national, protecting resources nationally can be rational.

A common deposit insurance fund changes this calculation. It bears losses across the group and, therefore, internalises both sides of an intragroup transfer. If moving resources from one subsidiary to another reduces total expected losses, there is no reason to block the transfer merely because the costs and benefits fall in different countries.

Our analysis, therefore, supports an important part of the current policy consensus. Common deposit insurance can remove the underlying incentive to ring-fence resources and allow capital to move to where it is most valuable.

The effects begin before a crisis

The implications do not stop there. If banks know that resources can be reallocated across subsidiaries in times of trouble, operating across borders becomes more valuable. A crossborder group has an internal capital market: a strong subsidiary can support a weak one, reducing the need to raise costly external capital. Ring-fencing reduces this benefit.

A common deposit insurance system, therefore, does more than just improve the functioning of existing crossborder groups. By making internal risk-sharing more effective, it also makes crossborder banking more attractive and encourages greater integration.

This is the logic that underpins much of the current policy debate: complete the missing pillar of banking union, and crossborder integration will follow. But there is a missing step. Better insurance changes behaviour.

When an impaired subsidiary can expect support from a healthy subsidiary, becoming impaired is less costly for the group. The prospect of costly external recapitalisation provides discipline; intragroup support weakens it. Better risk-sharing can therefore reduce incentives to avoid risk.

There is, however, an opposing force. A healthy subsidiary becomes more valuable because it can rescue another part of the group. When decisions are made at the group level, keeping one subsidiary healthy creates value not only there, but elsewhere in the group. This can strengthen incentives to avoid risk.

The effect of removing ring-fencing on risk-taking is therefore ambiguous. It depends on whether eliminating ring-fencing increases or decreases the value of the deposit insurance put.

When the fundamental risk in the economy is relatively low - that is, when subsidiaries are more likely to remain healthy for reasons unrelated to their own risk-taking - a troubled subsidiary is likely to find another subsidiary healthy enough to support it. The prospect of such support weakens the discipline imposed by costly external recapitalisation. Removing ring-fencing can therefore increase risk-taking.

When the fundamental risk in the economy is high, a healthy subsidiary is particularly valuable because it can preserve the value of the rest of the group. The group therefore has more to gain from keeping each subsidiary healthy, and removing ring-fencing can strengthen incentives to reduce risk.

The relevant policy question is therefore not simply whether common deposit insurance improves risk-sharing. It does. The question is how the prospect of better risk-sharing changes risk-taking.

More integration is not a sufficient welfare test

This distinction matters for welfare. A common deposit insurance system can reduce ring-fencing, improve resource allocation within banking groups, and encourage crossborder integration. But banks can take advantage of the insurance provided by internal capital markets by taking excessive risk, as they do not fully internalise the consequences of their risk-taking for the deposit insurance system. As a result, a reform can make crossborder expansion privately attractive even when the associated increase in risk imposes costs on the safety net.

When better intragroup risk-sharing strengthens incentives to keep subsidiaries healthy, risk-sharing improves, risk-taking falls, and common deposit insurance raises welfare. When it weakens discipline, however, the gains from moving resources more efficiently in a crisis must be weighed against greater risk-taking before the crisis. In our analysis, the latter can outweigh the former: common deposit insurance can encourage more crossborder banking while reducing welfare.

This creates a possibility that is largely absent from the policy debate. The policy implication is not to preserve national ring-fencing. But greater freedom to move capital and liquidity across borders should be accompanied by sufficiently strict supervision to address the risk-taking incentives created by expectations of future within-group support. ■

Gyöngyi Lóránth is Professor of Finance at the University of Vienna, Anatoli Segura is an Economist at the Financial Stability Directorate at the Bank of Italy, and Jing Zeng is Associate Professor of Finance at Universität Bonn

References

- Beck, T, J-P Krahnen, P Martin, FC Mayer, J Pisani-Ferry, T Tröger, B Weder di Mauro, N Veron and J Zettelmeyer (2022), [“Completing the Banking Union: Economic Requirements and Legal Conditions”](#), CEPR Policy Insight No. 119.
- Enria, A (2023), “Exogenous shocks and endogenous challenges: five years of European banking supervision (and beyond)”, Speech at the London School of Economics, 30 October.
- Financial Times (2026), “EU set to remove barriers to banks’ cross-border capital flows”, 19 June.
- Lóránth, G, A Segura and J Zeng (2026), [“Common Deposit Insurance, Cross-Border Banks and Welfare”](#), CEPR Discussion Paper No. 16893 (previous version circulated under the title “Voluntary Support and Ring-Fencing in Cross-border Banks”).
- Reuters (2026), “Europe antitrust chief urges EU countries to back cross-border bank deals”, 17 June.
- Schnabel, I and N Veron (2018), [“Breaking the stalemate on European deposit insurance”](#), VoxEU.org, 7 April.

Editor’s note: This column is based on CEPR Discussion Paper No. 16893 [“Common Deposit Insurance, Cross-Border Banks and Welfare”](#). This article was originally published on [VoxEU.org](#).



Reforming EU banking regulations

EU banking rules are under pressure for simplification. Jesper Berg and Hans Geeroms argue that reforms should cut complexity without weakening financial safeguards

Strong financial regulation provides benefits that often outweigh its costs. The financial sector requires close supervision because it channels private savings into the economy in a situation in which financial service providers understand finance much better than those they provide services to. Banks can also be so large ('systemic') that they pose a threat to financial stability, as seen nearly 20 years ago during the global financial crisis.

Nevertheless, the financial industry typically argues for less-strict regulation. Some European Union bank supervisors have joined the chorus, including the European Central Bank (ECB 2025, 2026), Germany's Federal Financial Supervisory Authority (*Bundesanstalt für Finanzdienstleistungsaufsicht*, Bafin) and central bank (Bundesbank 2025), authorities in Nordic countries¹ and the National Bank of Belgium (NBB 2026). Mario Draghi in his report on competitiveness for the European Commission said financial regulation had become too complex (Draghi 2024). The Dutch, German, French, Italian, Polish and Spanish finance ministers in March 2026 also called jointly for a simplification push².

In mid-July, the European Commission is due to publish a report on the competitiveness of the EU banking sector, which may take up these ideas³. To better understand the current debate on bank regulation, it is worth looking at the most interesting elements of three sets of proposals:

1. An ECB report published in December 2025 on possible simplification (ECB 2025a);
2. Bafin/Bundesbank (2025) on a simpler regime for small and non-complex banks;
3. October 2025 proposals from the European Banking Authority on the efficiency of the regulatory and supervisory framework (EBA 2025).

These proposals can be measured against a metric of whether they provide for simplification while not lowering regulatory standards. A secondary test is whether they would reduce complexity. The most promising proposal among those we review is a simpler regulatory regime, with safeguards, for smaller and non-complex banks.

Simplification proposals

The European Central Bank

A report written by an ECB High Level Task Force on Simplification (ECB 2025a) made numerous recommendations, most of which are not especially new and are unlikely to lower standards or have a substantial impact on simplification⁴. Many recommendations concentrate on regulatory and reporting simplification.

The financial sector requires close supervision because it channels private savings into the economy in a situation in which financial service providers understand finance much better than those they provide services to

However, three recommendations that touch on capital buffers that banks must hold as a safeguard are potentially concerning.

1. A proposal to reduce the number of capital elements in the prudential framework;
2. A recommendation that adjustments should be made to the design or roles of capital intended to absorb losses while a bank remains a going concern (known as Additional Tier 1 (AT1) capital) and capital to absorb losses once a bank fails (Tier 2);
3. Closer alignment of the EU Minimum Requirement for Own Funds and Eligible Liabilities (MREL) and the international Total Loss-Absorbing Capacity (TLAC) frameworks (standards for bank safety nets in case of financial crisis or bankruptcy), while reviewing how these interact with the going-concern framework. The EU introduced the stricter MREL standard in reforms in the mid-2010s, in advance of the entry into effect of TLAC in 2019.

These changes could lead to banks being required to hold less capital, potentially increasing the risk of bank failures. The capital requirements are admittedly complex. However, if elements are removed, the sum of capital requirements should not be reduced. Lowering capital requirements by removing certain elements would be simplification that lowers standards. Maintaining capital requirements while removing elements would be simplification without reducing standards.

AT1 mostly serves to generate profits for the investment banks that sell the assets that underpin it. AT1 has value for banks because it is uncertain, if a problem arises, whether authorities will impose losses on AT1 or resort to bail-in. This reduces the cost of AT1 to the banks but at a cost to taxpayers. AT1 should be abolished to create greater

certainty in resolution. It should be replaced by Common Equity Tier 1 (CET1) – the highest standard of capital. However, that would likely create an industry outcry.

The changes to align the MREL and TLAC requirements would be in line with some calls to reduce complexity by defaulting to TLAC, improve the efficient use of capital in Europe and reduce reliance on costly third-party markets⁵. This alignment is referred to as TLAC+. It is important that TLAC+ is not defined in a way that reduces capital requirements.

The Bafin/Bundesbank proposals

Bafin and the Bundesbank in a joint paper (Bafin/Bundesbank 2025) proposed a special regulatory regime for smaller and non-complex banks⁶. They have in mind the *Sparkassen* (savings banks) and *Genossenschaftsbanken* (cooperative banks), which dominate the German financial landscape, together with the two German globally systemic banks, Deutsche Bank and Commerzbank.

Smaller EU banks with total assets below €10 billion and simple business models, are subject to a disproportionate regulatory burden compared to their systemic risk, according to Bafin/Bundesbank (2025). These banks should be able to opt out of EU requirements⁷ that are based on the Basel standards for large international banks, including increasingly complex capital and liquidity requirements. They should be subject instead to a simpler, more proportional supervisory regime. This would replace risk-based capital standards with a leverage ratio, while streamlining reporting and stress-testing procedures.

A more proportionate regulatory framework for smaller EU banks (in the jargon: Less Significant Institutions, LSIs) can offer advantages. Nevertheless, history shows that multiple LSIs encountering difficulties simultaneously –

especially because of comparable underlying factors – may collectively exert a systemic effect on the financial sector. For example, about 25 smaller UK banks failed between 1991 and 1994 because of a crash in the domestic property market (Hindmoor and McConnell 2014).

Simpler standards for smaller banks should therefore be calibrated conservatively. Replacing risk-based capital requirements with a leverage ratio could otherwise encourage banks to increase their exposures to higher-risk assets.

Accordingly, the leverage ratio should be set above risk-weighted capital requirements and complemented by restrictions on activities such as securitisation and private-equity financing. Supervisors should also have powers to prevent regulatory arbitrage or inappropriate risk-taking, and to require banks to return to the full regulatory regime where necessary.

To preserve a level playing field, the largest small banks could opt into the larger-bank regime. If few do so, standards under the small-bank regime should be tightened.

With these safeguards, there is a strong case for not applying the strictest capital-buffer requirements to smaller banks.

European Banking Authority proposals

The EBA's recommendations cover, among other things, the production of technical standards and guidelines addressed to financial institutions – the EBA's core task⁸. In this area, the EBA proposes a new review methodology, a review of the existing rulebook and steps to make it more accessible (see Table 2 in the appendix).

The production of such technical standards and guidance suffers from at least three problems. First, such standards and guidance are produced when needed for the implementation of EU financial laws via secondary legislation. Issues that cannot be agreed in the top-level lawmaking process are pushed into secondary lawmaking (technical and implementing measures), recourse to which has exploded⁹. Many secondary laws go well beyond their initially more technical aim¹⁰.

Second, secondary acts are drafted by groups of specialists from all national authorities. They all have their own ways of doing things and the result is the sum of all approaches. Third, drafts then move through subcommittees and are finally approved by the EBA Board of Supervisors (BoS), the members of which are briefed by their specialists. Therefore, the BoS is rarely able to take a holistic view and make substantial changes.

To break this pattern, a change of culture is needed at all stages of the process. In this respect, the EBA has made the right proposals – but they must be delivered on. Peer reviews can strengthen mutual understanding of best practice, but legislative changes have made the process more adversarial between the EBA and the National Competent Authorities that work with the EBA¹¹.

Simplification should not spill over into lower standards

There is some justification for the view that EU financial regulation has become too complex. In this context, the ECB, Bundesbank/Bafin and EBA proposals – primarily on restructuring capital buffers and introducing a lighter-touch regime for smaller banks – deserve closer analysis. Proposals should be assessed according to two criteria: whether they make a difference, and whether they would lower standards.

A more proportionate framework for smaller, non-complex banks would be the most obvious win from a simplification push, provided it is calibrated conservatively and accompanied by robust safeguards against systemic

risk and regulatory arbitrage. Other simplification proposals are more questionable. The main danger is that the simplification debate will spill over into lower standards, particularly if a strong bank lobby uses simplification to push for lower capital standards.

Current capital and liquidity standards ensure well-capitalised banks and support long-term investment in Europe (Berg *et al* 2025). The global financial crisis cost the euro area about 9 percent of GDP (Chen *et al* 2019); that disaster should not be repeated. ■

Jesper Berg and Hans Geeroms are Non-Resident Fellows at Bruegel

Endnotes

1. See Swedish National Debt Office news of 9 May 2025, 'Debt Office in joint letter to EBA on rulebook for banking crisis management'.
2. The so-called E6. See 'E6 Letter - Advancing the Savings and Investments Union: An Imperative for European Competitiveness', 11 March 2026.
3. European Commission, 'Targeted consultation on the competitiveness of the EU banking sector', 11 February 2026.
4. Along with the High Level Task Force on Simplification report, the ECB published a report on banking supervision and simplifying macroprudential policy (ECB 2025b). See Table 1 in the appendix for further details of ECB recommendations.
5. See French Banking Federation, 'Improving the financing of the European economy: proposals of the FBF', 15 December 2025, and Villeroy de Galhau (2025).
6. See also Michael Theurer and Mark Branson, 'Europe's small banks need simplified rules', *Frankfurter Allgemeine Zeitung*, 12 September 2025.
7. The Capital Requirements Directive (CRD, Directive 2013/36/EU) and the Capital Requirements Regulation (CRR, Regulation (EU) No 575/2013).
8. In technical languages, guidance is provided in response to Level 2 and Level 3 mandates; see European Commission, 'Regulatory process in financial services', undated.
9. The most recent CRD/CRR update (see footnote 5) required more than 50 delegated acts. Anti-money laundering rules finalised in 2024 required close to 20 delegated acts.
10. A prominent example was the technical standards issued in 2014 laying out the definition of liquid assets for the purposes liquidity requirements, as required by Art. 460 CRR; see EBA press release of 28 March 2014, 'EBA publishes final draft Technical Standards on liquidity requirements'.
11. Including by reducing the role of the BoS and giving more power to EBA staff (as done in Regulation (EU) 2019/2175 of 18 December 2019).

References

- Bafin/Bundesbank (2025) 'A Simple Regulatory Regime for Small and Non-Complex EU Banks', unpublished non-paper, August.
- Berg, J, N Boivin and H Geeroms (2025) '[The quickly fading memory of why and when bank capital is important](#)', Working Paper 04/2025, Bruegel.
- Chen, W, M Mrkaic and M Nabar (2019) '[The Global Economic Recovery 10 Years after the 2008 Financial Crisis](#)', IMF Working Paper 2019/083, International Monetary Fund.
- Draghi, M (2024) [The future of European competitiveness](#), Publications Office of the European Union.
- EBA (2025) [Report on the Efficiency of the Regulatory and Supervisory Framework](#), European Banking Authority.
- ECB (2025a) [Simplification of the European prudential regulatory, supervisory and reporting framework](#), European Central Bank.
- ECB (2025b) [Streamlining supervision, safeguarding resilience: the ECB's agenda for more effective, efficient and risk-based European banking supervision](#), European Central Bank.
- ECB (2026) [Eurosystem response to the EU Commission's targeted consultation on the competitiveness of the EU banking sector](#), European Central Bank.
- Hindmoor, A and A McConnell (2014) '[Who saw it coming? The UK's great financial crisis](#)', Journal of Public Policy 35(1): 63-96.
- NBB (2026) [Annual Report 2025](#), National Bank of Belgium.
- Villeroy de Galhau, F (2025) '[Une approche européenne de la simplification : éviter trois idées fausses et proposer quelques jalons concrets](#)', Speech at Eurofi, Warsaw, 11 April.

This article is based on [Bruegel Analysis](#).

Appendix

Table 1. Proposals by the ECB high-level task force

Area	№	Recommendation	Specifics
Simplify the regulatory framework	1	Reduce the number of capital stack elements in the prudential framework.	Merge buffers into one releasable and one non-releasable. Keep P2G (Pillar 2 Guidance). Reduce number of buffer requirements.
	2	Improve the going-concern loss-absorbing capacity of the capital stack by adjusting the design or the role of AT1 (Additional Tier 1 capital) and Tier 2 instruments.	Either improve quality of AT1 or remove it.
	3	Expanding the degree of proportionality in the EU under the existing SNCI (Small and non-complex institutions) regime in a prudent manner.	Increase threshold for SNCIs and increase level of proportionality. The new simpler regime would be harmonised and calibrated in a more conservative manner.

	4	Reciprocate macro-prudential measures automatically up to a certain threshold, combined with, and conditional on, a more standardised application of macroprudential tools.	Hinges on more standard application of macroprudential tools. De minimis where not applied.
	5	Align the MREL (Minimum Requirement for own funds and Eligible Liabilities) and TLAC (Total Loss-Absorbing Capacity) frameworks more closely – without reducing going-concern resources – while reviewing their interactions with the going-concern framework.	TLAC floor and lower number of stacks on top.
	6	Refocus EU prudential law from directives to regulations, increasing harmonisation and regulatory transparency, and streamlining level 2 and 3 acts.	More regulations. Evaluate levels 2 and 3 to decrease prescriptive approaches and allow for increased supervisory discretion.

	7	Simplify the EU-wide stress test, streamlining the methodology and increase the usefulness of results from both a system-wide and bank-specific perspective.	Reforming stress tests, including considering top-down, to reduce compliance costs.
	8	Make the ECB Governing Council responsible for taking a holistic view of the overall level of capital demand within and across the banking union, while fully adhering to the principle of separation.	The Macroprudential Forum would perform a qualitative assessment of the adequacy and appropriateness of the overall level of capital requirements and buffers imposed on the SSM banking sector, abstracting from individual bank-level decisions.
	9	Finalise the savings and investment union, including completion of banking union, to reduce national fragmentation and allow for more efficient capital markets.	Concrete steps towards the finalisation of EDIS, with a clear timetable for implementation, and fostering deeper capital markets by progressing on the savings and investment union.

Simplify the supervisory framework	10	Strengthen and complete the Single Rulebook to simplify and harmonise supervisory practices and help deepen the internal market in banking.	Further harmonisation of rules and a review of national discretions.
	11	Reconsider the level of prescriptiveness of legislation governing supervisory processes to identify areas where more risk-based approaches can be implemented.	Give greater supervisory discretion in selected areas.
Simplify the reporting framework	12	Encourage European authorities to foster mutual data sharing, for example by operationalising the Better Data Sharing Regulation. Additionally, promote, through an EBA-led change management process, a regular coordination of EU-level data collections.	Less data requests through more coordination.

	13	Formulate a long-term vision within the JBRC (Joint Bank Reporting Committee), which would include monitoring the progress of initiatives to establish a fully integrated reporting system at European level for statistical, prudential and resolution purposes.	Eventually one database.
	14	Define a supervisory tolerance margin for minor errors to be disregarded.	EBA would be mandated to make a list of immaterial deviations or acceptable tolerances.

	15	Publish an inventory of non-market sensitive reporting requirements imposed on banks.	The European Commission should, either directly or by mandating the relevant authorities through level 1 regulation, ensure the publication of an inventory of all reporting requirements imposed on banks, including ad hoc requests outside of the established reporting frameworks. Step towards integrated reporting system. And to include all ESAs.
	16	Establish explicit obligations to conduct periodic assessments to validate the relevance and up-to-dateness of reporting requirements, based on objective criteria.	Conduct regular 'need to keep' reviews to eliminate unnecessary data.

	17	Fundamentally reform the EU public disclosure process by terminating the parallel transmission of supervisory and disclosure data.	Integrate supervisory and public disclosure reporting of pillar 3 data.
--	----	--	---

Source: Bruegel based on ECB (2025a).

Table 2. Selected proposals by the EBA for a more efficient regulatory and supervisory framework in the EU

Area	≠	Recommendation
1. Production of regulatory mandates	1	Set up a methodology to assess the materiality of L2 (Level 2) mandates.
	2	Review of the single rulebook by building blocks.
	3	Consolidated single rulebook
2. Reporting	4	Integrated reporting.
	5	Review and reduce existing reporting requirements.
3. Holistic picture	9	Reflect on the streamlining of capital/buffer/MDA (Maximum Distributable Amount) requirements and the multitude of tiers of own funds and TLAC/MREL.

	10	Introduce a more systematic application of simpler rules for the SNCI category while maintaining a single bank regime and explore the expansion of the SNCI category.
	11	Reflect on the existing balance of the home-host responsibilities and the use of waivers in the Single Market in the context of the Banking Union.
	16	Revisit the respective roles of L1/2/3, rely more on directly applicable rules.
4. Internal organisation	19	Rebalance EBA/CAs (national Competent Authorities) allocation of resources towards ex post convergence.
	20	Introduce top-down guidance to selected policy developments.
	21	Review the number of substructures, their mandate and strengthen their dialogue with BoS (Board of Supervisors).

Source: Bruegel based on EBA (2025).



Modernizing financial regulation

The financial system is constantly evolving, and the regulatory and supervisory framework must keep pace. Michelle Bowman discusses the work of the US and urges global financial watchdogs to be flexible

■ would like to discuss the work that is currently under way in the Financial Stability Board's Standing Committee on Supervisory and Regulatory Cooperation (or 'SRC') on modernizing financial regulation and supervision. Our efforts in the United States have been progressing swiftly over the past year, and most jurisdictions around the world have undertaken similar reviews of their respective regulatory and supervisory financial system frameworks.

Nearly a year ago, I assumed the chairmanship of the FSB's SRC. One of the SRC's workstreams under my leadership is the FSB's efforts to establish principles to guide regulatory and supervisory modernization around the world. Today, I will outline these principles, and I will share our approach to applying them in the United States.

The purpose of modernization and guiding principles

I will begin by describing modernization in this context. The financial system is constantly evolving, and our regulatory and supervisory framework must keep pace. We must ensure that our framework functions effectively and efficiently, preserves financial stability and bank safety and soundness, and promotes sustainable economic growth. Four principles should guide these efforts.

First, we must prioritize material financial risks—these are the risks that can impair an institution's viability and financial stability more broadly. Recent experience illustrates the danger of drifting from this principle. In the United States, we experienced the direct consequences of supervisors losing focus on material financial risks.

While we continue to understand the root causes that led to the failure of Silicon Valley Bank, evidence suggests that supervisors failed to sufficiently identify, understand, and act upon the material risks that ultimately contributed to the bank's failure.

At the time it failed, there were dozens of outstanding supervisory matters. But the sheer number of these findings clearly did not translate to supervisory effectiveness. The lesson is that more findings do not equal better supervision.

To implement effective modernization, we must move decisively away from a 'more is better' approach toward targeted, risk-based oversight. Therefore, we must become more disciplined in our approach to focus supervisory resources on what actually matters to an institution's safety and soundness.

We must ensure that our framework functions effectively and efficiently, preserves financial stability and bank safety and soundness, and promotes sustainable economic growth

Second, we must tailor regulation and supervision to be commensurate with each institution's risk profile. Banks differ dramatically in size, systemic impact, complexity, and risk profile. A community bank offering only traditional lending products does not present the same risks as a multi-trillion-dollar institution with global operations and complex trading activities. When we design standards for the most complex institutions and apply them to simple, non-complex institutions, we create inefficiency and burden without corresponding benefits or improved outcomes.

Third, we must prioritize transparency and accountability in our regulatory and supervisory processes. Regulatory and supervisory expectations should not come as a surprise. Public disclosure of our supervisory practices and expectations for regulated firms enables supervised institutions to understand what supervisors expect and how supervision is conducted.

This allows both firms and supervisors to better manage risks and allocate resources more effectively. Transparency also strengthens accountability and ensures that all parties—supervisors, regulated institutions, boards of directors, and the public—share a common understanding of supervisory objectives and methodologies. This clarity ensures that prudential oversight is effective and efficient.

And fourth, modernization must be forward-looking so that it considers emerging risks and promotes responsible innovation. Rapidly evolving technologies offer significant benefits and present complex challenges for the financial sector.

To maintain a robust banking system, supervisors must support responsible innovation while continuing to address material risks. That requires regulatory and supervisory frameworks that are not overly prescriptive, and which draw upon real-world experience to establish principles that can encompass specific circumstances and use cases.

Modernization in practice

The FSB has taken stock of modernization activities across its membership, including in standard-setting bodies. These modernization efforts are widespread and emphasize similar objectives, including keeping frameworks forward-looking and adapting to current and emerging material risks, enhancing tailoring and proportionality, reducing unnecessary burden, enabling responsible innovation, using data and technology-enabled supervision, and supporting economic growth.

The next stage of this work is a report for public consultation that will be published in the fall. It will include principles to help guide the implementation of FSB member modernization activities. The report will then be delivered to the G20 as an important milestone in our shared commitment to modernization.

Over the past year, the United States has made significant progress on our work to modernize both the US regulatory and supervisory frameworks. These efforts demonstrate how the principles work in practice.

Capital framework reforms

One of the primary US modernization priorities includes reforming the capital framework. The Federal Reserve's 2026 Basel III proposal demonstrates how modernization principles translate into concrete regulatory improvements that balance sound prudential standards with economic efficiency.

In March, we published comprehensive proposals to reform the US capital framework for public comment. Rather than working backward from a predetermined aggregate capital target, we evaluated each requirement on its merits, examining whether it is appropriately calibrated to risk, achieves its intended purpose, and minimizes unintended consequences.

The proposed capital framework will simplify requirements in several ways: by moving large banks to a single stack of risk-based capital requirements, adopting a true risk-based capital approach; recalibrating the global systemically important banks (G-SIB) surcharge to better reflect systemic risk; reducing overlaps between stress testing and risk-based requirements; and right-sizing requirements to reduce incentives for traditional banking activities to migrate outside of the regulated banking sector. The proposal also commits to indexing the G-SIB surcharge to nominal economic growth going forward. This will fulfil a decade-old commitment to prevent unintentionally increasing requirements¹.

The comment period recently closed, and we are now evaluating public feedback and working to finalize these rules. The goal is to maintain robust capital levels while reducing regulatory barriers that constrain credit availability without providing corresponding benefits. This bottom-up, risk-based approach illustrates how modernization can enhance both effectiveness and economic efficiency.

Statement of supervisory operating principles

We recently enhanced our supervisory framework, starting with increased transparency and accountability by publishing Supervisory Operating Principles to guide the activities and priorities of our supervision. This is the first time the Federal Reserve has disclosed the principles that form the foundation of our supervisory approach.

The focus on early detection and mitigation of material risks provides clarity for supervised institutions and the public about how we conduct supervision. As we refine our practical implementation, we will continue to evaluate and enhance these Principles.

Updating asset thresholds

As with any aged infrastructure, many of the requirements contained within our regulations and supervision are not

aligned with the growth of our economy. Therefore, we are working to address long outdated fixed asset thresholds that ignore economic growth and inflation. A bank with \$10 billion in assets today is not the same as a bank with \$10 billion in assets a decade ago, and banks with unchanged risk profiles face increasing regulatory burden over time, simply due to economic growth and stagnant thresholds. In addressing this issue, we are working to craft solutions that will grow over time by indexing fixed dollar thresholds.

Refining material financial risk assessment

Implementing the operating principles to prioritize financial risk management requires refining our approach to assessing material financial risks. In doing so, we expect our supervisors to exercise reasoned judgment throughout the examination process, from pre-exam planning through the exam exit process. Our supervisory programs continue to require comprehensive supervision, but our findings should focus on those issues that impact an institution's safety and soundness.

In practice, this means revised standards for our examination findings that incorporate objective, measurable factors to evaluate material financial risks. Those matters that raise concern, but do not reach the severity level of a violation can be cited as 'supervisory observations' for notable, but less significant issues.

Emerging technology

Given the rapidly evolving state of technology, the US has expanded modernization efforts to include allowing banks to safely adopt emerging technologies. Last month, under a second FSB SRC workstream, we published a consultation report on sound practices for financial institutions as they consider adopting and using artificial intelligence (AI)². The adoption and integration of AI in the financial system offers significant benefits, but its rapid evolution presents unique challenges.

The report draws on real-world examples of how banks and other financial institutions are already using AI successfully and responsibly. Rather than prescribing one-size-fits-all requirements, it offers practical guidance that institutions can adapt to their specific circumstances and use cases.

The FSB is seeking public comment on the report through July 22, and your feedback will ensure this work is useful for the industry. Public feedback will also assist our efforts to apply the same principles of flexibility, risk-focus, and practicality that guide all of our modernization efforts.

Additional US modernization work has improved supervisory guidance, simplified small bank capital frameworks, and enhanced supervisory stress testing.

Going forward

The FSB has a role in fostering international cooperation and sharing best practices. To ensure the effectiveness of our work together, the FSB should recognize the principles I have outlined today as each jurisdiction implements ongoing modernization efforts.

At the same time, the FSB should encourage flexibility among approaches to ensure that modernization efforts are appropriate for individual jurisdictions. Efforts to prescribe and enforce strict rules that are not suitable for the diversity of institutions and authorities within individual jurisdictions ultimately erodes the FSB's effectiveness.

In the United States, we are committed to enacting these modernization principles and welcome similar efforts. We stand ready to share our experiences, learn from others, and work together toward common goals of financial stability, innovation, and economic growth.

Modernization is an evergreen approach to ensuring financial stability—starting with the banking sector. It is an ongoing commitment to improve, to learn, and to adapt our approach as markets, technology, and risks evolve.

By following the principles and examples outlined in my remarks, the FSB can establish a strong foundation for regulatory and supervisory frameworks that learn from the past, address the present, and prepare for the future. ■

Michelle W Bowman is the Vice Chair for Supervision of the Board of Governors of the US Federal Reserve System

Endnotes

1. See Board of Governors of the Federal Reserve System, “Regulatory Capital Rules: Implementation of Risk-Based Capital Surcharges for Global Systemically Important Bank Holding Companies,” 80 FR 49082 (August 14, 2015) (“To ensure changes in economic growth do not unduly affect firms’ systemic risk scores, the Board will periodically review the coefficients and make adjustments as appropriate”).
2. The report is available on the FSB’s [website](#).

The views expressed here are my own and are not necessarily those of my colleagues on the Federal Reserve Board or the Federal Open Market Committee. This article is based on a [speech](#) delivered at a Bank Policy Institute London Conference, London, 13 July 2026.

A strategic lever to boost global decarbonisation

The Paris Agreement provides a framework for international carbon trading. Georg Zachmann and Annabelle Weisser and argue that international carbon credits can support the EU's 2040 target if buyers' club rules and leverage ensure real mitigation gains abroad

Executive summary

Article 6 of the Paris Agreement provides a framework for international carbon trading under which countries can agree to buy and sell emissions reductions to meet their national climate targets, or for other international mitigation purposes, such as in aviation. In this context, the European Union will allow up to five percentage points of its 2040 climate target to be met through purchases of international credits.

This opportunity should be used not merely as a cost-saving mechanism but should be a tool to encourage and increase global emissions reductions. We propose a combination of two institutional innovations: a Carbon Buyers' Club that would ensure common quality standards, and a set leverage ratio to ensure that every unit of domestic emissions offset leads to a greater amount of mitigation abroad.

Within a market-based architecture, private buyers should be able to use international credits as collateral to borrow domestic allowances. Through this controlled exchange against EU allowances, the system would link international flexibility to domestic compliance systems, in this case the EU emissions trading system.

To safeguard environmental integrity our proposal would move away from project-based crediting. Instead, partner countries would allocate portions of their own emissions reductions to other countries, and would receive in return direct financing, project-based investments or both.

1 Introduction

Compared to much of the rest of the world, Europe has had some success in cutting emissions. The European Union's share of global greenhouse-gas emissions dropped from 15 percent in 1990 to 6 percent in 2024¹. However, even if the EU were to reach net zero tomorrow, the total impact on global warming would be limited.

To meaningfully contribute to global emissions reduction efforts, in a context in which the world is far from keeping global warming well below two degrees Celsius above pre-industrial levels (UNEP 2025), the EU must leverage its impact not only domestically but also abroad.


To keep climate protection alive, the EU should focus its climate policy on its global impact and design its international credits procurement strategy accordingly

The EU's first tool to push for more climate change mitigation internationally is diplomacy. EU diplomacy was a major driver of the United Nations Framework Convention on Climate Change (UNFCCC) process that culminated in the 2015 Paris Agreement, which provides a universally accepted legal basis for policy instruments that the EU and others can use to catalyse emission reductions beyond their borders. The UNFCCC process is thus also a major driver of EU climate policy (Pisani-Ferry *et al* 2023).

Beyond its diplomatic weight, the EU has used five main levers to encourage climate change mitigation abroad (Figure 1):

1. By pledging ambitious domestic emission reductions (Paris Agreement, Art. 4), the EU sets an example for others to also make significant mitigation commitments. To meet its ambitious targets, the EU has developed policies, technologies and expertise, which in principle lower costs and reduce uncertainties for countries that follow suit².
2. The EU is the world's largest provider of climate finance (Art. 9)³. This support encourages developing countries to make climate commitments and reduces their capital costs for mitigation investments.
3. By facilitating technology transfer (Art. 10), the EU could in principle enable developing countries to gain more affordable access to mitigation technologies.
4. By implementing green trade mechanisms (such as the carbon border adjustment mechanism, CBAM), the EU can encourage trade partners to introduce mitigation policies and decarbonise their industries.
5. International carbon markets (Art. 6) can encourage substantial emission reductions in countries with low mitigation costs.

Figure 1. International carbon trading as a pillar of EU international climate policy



Source: Bruegel.

However, some levers have more potential than others. There is limited scope for the EU to further tighten domestic climate targets or to provide dozens of billions of euros in additional international climate finance. The EU's diplomatic weight and technology leadership have not increased since the Paris Agreement was negotiated. But green trade policies are a promising development (Bahí *et al* 2026), and international carbon trading is making a return.

A framework for international carbon trading is provided by Article 6 of the Paris Agreement. This allows countries to cooperate bilaterally by transferring emission reductions, known as Internationally Transferred Mitigation Outcomes (Article 6.2). Buyers can use these transfers to meet their national climate targets or for other international mitigation purposes, such as in aviation. To maintain environmental integrity, these transfers must be carefully accounted for and tracked throughout their lifecycle using reliable national or international registries.

Article 6 also established a centralised, UN-supervised international carbon crediting mechanism (Article 6.4). Under this framework, approved mitigation activities generate unique emission reduction credits that are recorded and traced within a central registry.

Finally, it provides a framework for non-market cooperation by countries to reduce emissions (Article 6.8), which does not involve the generation or transfer of tradable carbon credits.

If well designed, international carbon trading can become a strategic instrument in the EU's toolbox to encourage international climate change mitigation. At the climate summit (COP29) in Baku in 2024, UNFCCC participants finalised the Article 6 rulebook, allowing countries to procure and use international carbon credits⁴ to achieve climate targets.

In 2025, the EU decided to allow up to five percentage points of its 2040 climate target – under which emissions should be cut by 90 percent compared to 1990 – to be met by high-quality international credits (Regulation (EU) 2026/667).

Such trade in mitigation outcomes can lower the global cost of decarbonisation, as substantial volumes of abatement, especially in emerging market and developing economies (EMDEs), are still relatively cheap⁵.

However, there are considerable risks involved in international carbon trading. Previous experiences exposed profound weaknesses. Under the Kyoto Protocol's Clean Development Mechanism (CDM), many credits were issued for claimed emission reductions that were likely overestimated and may have happened regardless of whether or not credits were issued (Cames *et al* 2016). These concerns were the motivation behind the San José Principles, which advocated stricter integrity rules⁶.

However, the safeguards ultimately adopted in the Article 6 rulebook fell short of these standards⁷. This failure to implement higher-integrity rules undermines trust in international carbon markets and fuels justified scepticism towards offsetting (Asadnabizadeh and Moe 2024).

If credits on paper count towards buyer countries' climate commitments without corresponding to real emission reductions in seller countries, they risk undermining the fundamental architecture of the Paris Agreement. And without sufficient safeguards, Article 6 risks creating a race to the bottom: buyer countries are incentivised to minimise compliance costs, while seller countries seek to maximise credit selling. The fundamental challenge is to create incentives for all participants to adhere to environmental integrity standards.

These integrity concerns are particularly pronounced when credits are sold by individual project developers. Since project-based credits rely on hypothetical estimates of what emissions would have been without the project, it is often difficult to verify whether claimed emission reductions are real, additional and permanent (Fuchs *et al* 2026). While safeguards in the Article 6 rulebook are intended to improve the integrity of credits, they cannot eliminate the fundamental uncertainty around project-based credits.

Despite the challenges, the EU has limited time to set up a functioning carbon market. Its five-percentage point allowance for emissions cuts to be met through 'high-quality international credits' (Regulation (EU) 2026/667) applies from 2036 (Montrone *et al* 2026). A pilot period from 2031 to 2035 will establish a market for these credits.

Strategic decisions on how international carbon credits will be used in the EU need to be taken now and are expected in 2026/2027. In the framework of a review of the EU emissions trading system (ETS)⁸, the European Commission proposed on 17 July 2026 to set up a separate compliance mechanism for international credits (European Commission 2026).

However, the Commission proposal is only the beginning of the conversation, and the Commission is expected to present a *"framework for the possible use of international credits towards the 2040 EU climate target under the European Climate Law"* in late 2026⁹. Together with other legal acts, the EU needs to develop a consistent climate policy framework after 2030 that includes the strategic role of international credits¹⁰.

This *Policy Brief* outlines a proposal for how to make Article 6 work to catalyse global emission reductions. The proposal rests on three principles: i) country-issued credits instead of project-based ones, ii) private buyers instead of public buyers, iii) and a leverage ratio that implies that the EU should use international carbon trading to encourage mitigation abroad, rather than just transferring carbon emissions from one region to another.

Section 2 spells out these principles as fundamental choices, including arguments for why certain avenues should be prioritised over others. Section 3 outlines how this could work in practice. Section 4 discusses conceptual uncertainties.

2 Fundamental choices

Strategic questions about how the EU uses Article 6 include who procures foreign credits (the EU, national governments or companies), in which compliance mechanisms they will be accounted for and how to determine the minimum acceptable quality of foreign credits (Delbeke and Bencini 2026).

2.1 Who sells credits? Project developers vs. governments

We recommend moving away from project-based mitigation trading to issuance of carbon credits at national level. The certification of individual mitigation projects poses several environmental-integrity risks related to the measurement, questionable permanence and additionality of credits (Fuchs *et al* 2026; Edenhofer *et al* 2026).

These challenges make it nearly impossible to assess consistently how much additional mitigation projects deliver over time. In contrast, countries could make parts of their own emissions reductions available to count towards other countries' climate commitments – the Nationally Determined Contributions (NDCs) (Delbeke and Bencini 2026). Seller countries would retain full discretion over whether they receive direct financing, project-based investments or a combination of both in return.

This builds on existing concepts: the international transfer of carbon credits under Article 6 triggers a 'corresponding adjustment' in the host country's NDC; ie. by buying foreign credits, buyers already obtain a share of the seller country's carbon budget.

At the same time, host countries can require ‘benefit sharing’ from project developers in the form of a share of credits or a tax. In our proposal, seller countries, rather than project developers, are the issuers of carbon credits. This implies that the value of a credit depends not on the mitigation achieved by a specific project, but on a country’s mitigation success compared to a baseline emission pathway.

2.2 Who buys credits? Governments vs. private companies

As an alternative to public entities buying foreign credits, private companies should buy foreign mitigation outcomes to exchange at home for domestic allowances. If public entities such as the European Commission or national governments buy foreign credits, there is a risk of politicising the trade.

To domestic audiences, the purchase of foreign credits with public money could look like substantial hand-outs to companies and foreign governments. Internationally, countries unable to sell because of a lack of mitigation capacity might perceive this as a political rather than environmental or economic choice.

However, if private companies can buy foreign credits to obtain domestic compliance relief, they will try to get the highest volume of domestic compliance at the lowest price. If the domestic compliance value of a specific seller country’s credit can be made to align with the effective emission mitigation of that country within a certain timeframe (eg. 2030-2040), private companies will have an incentive to deploy market-based information systems to determine the value of each seller country’s credits.

In the EU, private companies would have an incentive to buy foreign mitigation outcomes if they can be used for compliance in the ETS or the separate ETS2, which covers emissions related to buildings and road transport¹¹.

2.3 Offset ratio: one-for-one vs. leverage

Finally, the purchase of foreign credits should be designed to accelerate global decarbonisation. The EU simply replacing one unit of domestic mitigation with one unit of foreign mitigation will not drive global decarbonisation.

Instead, the trade in foreign mitigation credits should reflect that in some countries, mitigation is cheaper than in the EU. With a policy-set leverage ratio of 1:5, for instance, every tonne of emitted carbon dioxide that the EU offsets could correspond to five tonnes of emission reductions abroad¹². This can be seen as a cautious implementation of Article 6.4, which demands that international transfers entail additional emission reductions¹³.

By shifting from one-to-one offsetting to a high-ambition multiplier, developed countries can turn their need for domestic flexibility into a major driver of global climate action. If a coalition of countries in the Organisation for Economic Co-Operation and Development (OECD), excluding the United States¹⁴, capped its use of international credits at five percent of 1990 emission levels¹⁵ during the final 15-year transition to net zero (2036-2050)¹⁶, it would offset less than seven gigatonnes of hard-to-abate emissions in those countries.

By requiring a 1:5 mitigation ratio for these purchases, this coalition would stimulate nearly 34 gigatonnes of mitigation¹⁷ in partner countries. Even after allowing for slightly higher emissions in buyer countries, this strategy would yield a net global emission reduction of 27 gigatonnes¹⁸.

With the world currently on track to overshoot the two-degree carbon budget by maybe 125 gigatonnes¹⁹, implementing this high-leverage procurement policy would allow OECD buyers to close more than a fifth of the entire global emissions gap. Table 1 summarises the three fundamental choices.

Table 1. Fundamental choices in international carbon market design

Strategic question	Alternatives	Proposed approach	Main rationale
Who sells mitigation outcomes?	Project developers vs. governments	Governments	Tackles integrity and additionality issues and creates stronger country-level incentives
Who buys mitigation outcomes?	Governments vs. private companies	Private companies	Facilitates market-based pricing mechanisms, mobilises private finance and avoid politicisation
What is the offset ratio?	One-for-one vs. leverage	Leverage	Ensures that offsetting generates additional mitigation

Source: Bruegel.

3 Putting the framework into practice

To ensure the leverage effect of international carbon trading, the system would rely on three mechanisms:

1. A mitigation value that would reflect the actual emission reductions a seller country achieves relative to its baseline;

2. A Carbon Credits Clearing House that would allow companies to borrow domestic compliance allowances by collateralising adequate amounts of foreign credits;
3. A Carbon Buyers' Club that would coordinate high compliance standards among participating countries.

3.1 Mitigation value

The mitigation value would quantify a foreign credit based on the volume of emissions reductions it represents beyond what the seller country would have achieved in the absence of the trade. This metric would enable direct comparisons between different seller countries, using an EU allowance with a value of one as the standard benchmark.

At the core of the mitigation value concept is a baseline of emissions reductions that a country should achieve to keep global warming well below two degrees (on the corresponding challenges, see section 4). A country's emissions reductions that go beyond this baseline can be sold. The value of a credit depends on the extent to which the baseline has been overachieved and the quantity of credits that the country has issued:

$$\text{Mitigation value} = \frac{\text{overachievement of baseline}}{\text{international allowances sold}}$$

The mitigation value of a foreign credit would determine the volume of domestic emissions it can offset. The definitive mitigation value of a unit will only be known ex post, in 2040, because it will depend on how many credits a country sells over time and the extent to which the country actually decarbonises.

To get the best value for money, companies wanting to buy credits up to 2040 have an incentive to investigate whether individual seller countries are successful in mitigating emissions compared to the volumes of credits they are issuing. The market will thus develop its own dynamic ex-ante evaluation of the mitigation values for seller countries' units.

While the direct exchange of foreign credits into compliance units only happens once the mitigation value of foreign credits is established in 2040, market participants will want to use foreign credits for domestic compliance earlier. A financial institution is needed to bridge this gap: a Carbon Credits Clearing House.

3.2 Carbon Credits Clearing House

A Carbon Credits Clearing House (CCCH) would operate as an exchange platform to bridge the gap between purchases of foreign credits and the moment their definitive values are known. A company would buy foreign credits and deposits them in the CCCH. It would then use these credits as collateral to borrow domestic allowances from the CCCH, which it can use immediately for domestic compliance.

The overall quantity of EU allowances available at the CCCH must not exceed five percent of 1990 emissions, reflecting the limit in the EU's 2040 climate target.

The CCCH would determine how many domestic allowances a company can borrow against the number of foreign credits it deposits. In doing so, the CCCH would discount the value of foreign credits used as collateral based on their expected country-specific mitigation value.

To achieve this, the CCCH will monitor the market and align its own collateral requirements with the market's assessment of the mitigation values assigned to individual countries.

Moreover, the CCCH's collateral requirements would take into account the policy-set general leverage ratio – for example, 1:5. For instance, for a collateral of 1,000 credits from a foreign country, the CCCH might issue only 100 EU allowances if the agreed leverage rate is 1:5 and the country's expected mitigation value is 0.5²⁰.

The final mitigation value will be determined in 2040 by the Carbon Buyers' Club (see section 3.3). For example, if in 2040 it turns out that the country from the example above has achieved 0.4 gigatonnes (Gt) of emissions reductions beyond its baseline and has sold allowances equivalent to 0.5 Gt, its mitigation value would be 0.8 (0.4 Gt/0.5 Gt).

The mitigation value of the country's credits would be multiplied by the policy-set leverage ratio of 1:5. One EU allowance would thus correspond to 6.25 of that country's credits ($0.8 \times 1:5 = 1:6.25$). For a company that initially borrowed 100 EU allowances against a collateral of 1,000 foreign credits, the 1,000 foreign credits would then be worth 160 EU allowances ($1,000/6.25$) and the company would receive 60 additional EU allowances. Figures 2 and 3 illustrate this mechanism.

3.3 Carbon Buyers' Club

Catalysing global mitigation requires collaboration with other committed countries to avoid substandard trading and to pool leverage. So far, our proposal has addressed how carbon trading could be designed between the EU and those countries that want to make available portions of their national mitigation capacities.

To increase its impact, the EU should engage with partner countries – in the OECD, for example. Some OECD partner countries, such as Switzerland and Japan²¹, have already begun using international carbon credits under Article 6.2 of the Paris Agreement²².

A Carbon Buyers' Club could coordinate among those countries that use international credits to meet their own climate targets. It would set common standards for counting international allowances toward domestic climate targets.

In our proposal, the Carbon Buyers' Club could be involved in determining the baseline for each potential seller country, which would determine the threshold above which emission reductions could be made available for purchase. It would also set out the methodology and possibly the governance system that conducts the final evaluation of each seller country's mitigation value (reflected in Figure 3).

The Carbon Buyers' Club could also try to coordinate buyers' leverage rates to encourage the highest possible global mitigation arising from international carbon trading. As discussed in Zachmann (2025), a coalition of buyers with sufficient leverage rates could make a substantial contribution to push global emissions closer to keeping global warming well below two degrees.


However, the club should not serve to create unfair buying power in terms of pricing. Each buyer (companies or governments, depending on the model each country chooses) would still compete to buy credits with high expected mitigation values.

4 Conceptual challenges

Our proposal if taken up would mark a significant step away from the international carbon markets that exist in practice and the concepts that have dominated the discussion since the Baku climate summit (COP29). Our proposal suggests three strategic choices:


1. Shifting the issuance of carbon credits away from individual project developers and towards governments;

Figure 2. Using foreign units as collateral, the CCCH lends domestic compliance units



Note: EUA = EU allowance; each EUA is equal to one tonne of carbon dioxide equivalent (CO₂e) within the ETS.
Source: Bruegel.

Figure 3. After stocktake evaluation in 2040, private buyers settle their accounts



Source: Bruegel.

2. Shifting the purchase of foreign units away from governments and towards private buyers;
3. Applying a leverage ratio that recognises international carbon trading as a strategic lever to encourage greater global mitigation, instead of merely shifting emission reductions from one region to another.

These innovations serve different purposes: first, carbon credits issued at country level could solve many of the integrity, additionality and permanence issues that dog project-based units²³. Rather than focusing on hypothetical estimates of emissions that would have existed in the absence of a project, country-level credits would cover emission reductions that will actually be achieved by 2040. Second, private buyers purchasing units could avoid politicisation, efficiently process market information and leverage private finance. Third, a leverage rate would ensure that international transfers result in additional emission reductions, rather than just shifting one unit of mitigation from one region to another.

However, the outlined proposal is not a complete blueprint. Solutions to important political, legal, institutional and technical challenges still need to be found. The European Commission should play a central role in developing a mechanism that will use international flexibilities to encourage substantial additional emission reductions. Based on an ambitious framework, it should explore and negotiate with the main domestic and international stakeholders to find a balanced solution.

The most politically difficult point will be setting baselines for seller countries' emissions against which improvements are to be sold. Setting the overall budget (broadly implying a maximum warming limit) and attributing it across countries is not a technical question; it is highly political.

A well-below two-degree warming target in line with the Paris Agreement could inform the overall budget. Allocating it could be done based on an indicator that combines different criteria²⁴, ensuring that for each country this compound indicator defines a country-specific baseline that is more ambitious than its current NDC. However, the exact point at which the baseline is set implies decisions on how the financial burden of mitigation should be shared.

This requires a comprehensive assessment of different perspectives and fairness criteria. Setting the baseline also requires careful calibration of the trade-off between ambition and incentives. A too-low baseline risks oversupplying credits and lowering global ambition, while overly ambitious baselines may reduce incentives for countries to engage in carbon trading in the first place²⁵.

A related challenge is how to reflect the impact of leveraged carbon trading in countries' emissions balances. The proposal implies that emissions in seller countries would decline much more than what they offset in buyer countries. Accounting for this difference in countries' emissions balances, in the context of the Paris Agreement architecture, requires careful consideration.

Another basic legal question is how to make the proposal compatible with the Paris Agreement and its institutions without requiring unanimous support by all parties to the Paris Agreement. The governance of the Carbon Buyers' Club and the institutions that assess the final mitigation values and operate the collateralisation and issuance of EU allowances will require thoughtful design.

The proposed architecture also raises institutional and technical questions about the finite time horizon over which the mitigation value of international credits can be assessed. In particular, the allocation of risk requires careful consideration.

If the mitigation value of foreign credits proves lower than expected in 2040, purchasing companies should bear the primary loss, with the CCCH only exposed if firms become insolvent. If domestic carbon prices fall, firms that invested in more expensive international credits may incur losses as well.

On the other hand, if the mitigation value of foreign credits is higher than expected, firms receive additional EU allowances that they can sell or use for domestic compliance (see example in section 3). The system inherently entails a risk for buyers, which policymakers could help mitigate by designing hedging instruments, such as carbon contracts for difference (McWilliams and Zachmann 2021).

Setting up a functioning international carbon market would also benefit from more clarity about the post-2040 period. EU legislation foresees the use of Article 6-type flexibilities only from 2036 to 2040. If this time limitation is maintained, the market will lack the long-term certainty required to develop robust, functional infrastructure.

Conversely, if exchanges of credits successfully support EU policy aims, they might be useful beyond 2040 (as illustrated in section 2.3). Policymakers should already outline how international carbon trading could evolve into a permanent feature of the EU climate policy framework, providing the long-term predictability needed for the market to develop.

These are only some examples of the many legal, political, diplomatic, economic questions that will arise in the preparation, design and implementation of the presented proposal.

5 Conclusion and outlook

The EU's future use of international carbon credits presents a strategic choice. It can either focus on lowering the cost of domestic compliance, or it can design a tool to strengthen global decarbonisation efforts. But without a pathway towards global emissions reductions, aggressive domestic mitigation risks becoming politically untenable.

To keep climate protection alive, the EU should focus its climate policy on its global impact and design its international credits procurement strategy accordingly.

Achieving this objective requires rethinking how international carbon markets can be organised. Three main conceptual innovations could prove fruitful:

1. Tradable mitigation outcomes should be issued by governments rather than individual projects, linking transfers directly to national mitigation performance and reducing many of the integrity-related challenges associated with project-based crediting.
2. Private companies, rather than governments or EU bodies, should purchase these credits. This would avoid the risk of politicisation, harness market incentives to identify high-quality credits and mobilise private finance for global decarbonisation.
3. The system should include a government-set leverage ratio, ensuring that international flexibility contributes to overall global mitigation rather than merely redistributing existing efforts.

Implementing such a system requires an institutional foundation. Mitigation values linked to countries' actual performances would provide the basis for valuing international credits. A Carbon Credits Clearing House could bridge the time gap between purchases of foreign credits and final assessments of their mitigation value, allowing companies to use international units for domestic compliance before their full value is known.

Ensuring that Article 6 does not become a loophole for rich countries, but is an instrument to reinvigorate global decarbonisation, relies on the EU not acting in isolation. A multilateral Carbon Buyers' Club would help facilitate

the required coordination. This coalition would empower participating countries to collectively harmonise quality standards, methodologies and leverage ratios, thus safeguarding market integrity against a race to the bottom.

By enabling productive use of Article 6 by a large coalition of buyer countries, the Carbon Buyers' Club would ensure the integrity of traded credits, thus stabilising the fundamental architecture of the Paris Agreement.

Several political, legal and technical questions remain unresolved, including the determination of country baselines, the governance of the proposed institutions and the finite time horizon. These challenges are substantial but should not obstruct harnessing the broader opportunity.

The EU has about ten years to set up the mechanism and negotiate it with both seller and buyer countries before international carbon credits begin counting towards its emission reductions target in 2036. This period should be used to develop a system that treats Article 6 not simply as a cost-containment instrument, but as an important lever for encouraging additional mitigation efforts. ■

Georg Zachmann is a Senior Fellow and Annabelle Weisser is an Energy and Climate Programme Officer, both at Bruegel

Endnotes

1. EU emissions dropped from 4.9 gigatonnes of CO₂ equivalent in 1990 to 3.2 Gt in 2024, while global emissions increased from 32.2 Gt to 53.2 Gt over the same period (Crippa et al 2025).
2. One example of how the EU has created a mitigation infrastructure is the generous deployment subsidies for renewable energy, which allowed the sector to scale up and contributed to rapid cost declines (Nicolini and Tavoni 2017).
3. See the Council of the EU explainer of 9 January 2026, [‘Europe’s contribution to climate finance \(in € billion\)’](#).
4. There is no common terminology for the traded items. For sake of simplicity, we use ‘credit’ for certificates of emission reductions that can be traded across countries. We use ‘allowances’ for EU allowances in the EU ETS, where one allowance corresponds to one tonne of carbon emissions; see European Commission, [‘About the EU ETS’](#).
5. Reasons for higher marginal abatement costs in developed economies include path dependencies and higher land and labour costs (Glennster and Jayachandran 2023).
6. For instance, the San José Principles advocated for an Article 6 rulebook that “prohibits the use of pre-2020 [CDM] credits”. See the website of the Government of Costa Rica, [‘About the San José Principles’](#).
7. By comparison, the finalised Article 6 rulebook permits the use of some pre-2020 CDM credits (Decision 20/CMA.7).
8. The 2023 revision of the EU ETS Directive requires an assessment on shaping the carbon market’s legislative framework beyond 2030 and aligning it with the EU’s 2040 climate target.
9. See European Commission, [‘Call for evidence for an impact assessment: European Climate Law - Climate Action’](#).
10. See European Parliament, [‘Revision of national targets and flexibilities in the EU climate policy framework’](#).
11. See European Commission, [‘ETS2, Buildings, road transport and additional sectors’](#).
12. With one-for-one offsetting, the arbitrage gain is shared between buyers and sellers – buyers because of lower compliance costs and sellers because of additional income. The idea of the leverage ratio implies that the price difference of mitigation outcomes can be utilised to benefit the climate by facilitating additional emission reductions. A political economy challenge is how exactly to set the rate to maximise global emission reductions while still allowing for some reduction in compliance cost.

13. The Paris Agreement Crediting Mechanism (PACM) aims “to deliver an overall mitigation in global emissions” (Paris Agreement FCCC/CP/2015/10/Add.1, Article 6.4).
14. The United States is excluded from this calculation because of its withdrawal from the Paris Agreement, which took effect in January 2026 (CRS 2025).
15. Non-US OECD countries’ emissions were 9.0 Gt. See OECD Data Explorer, [‘Air emissions - Greenhouse gas emissions Inventories’](#).
16. We use 2036 as the starting point for this calculation because the EU will start using international credits for its own climate targets then (see section 1). We expand the calculation to 2050 to show the potential of the mechanism, despite the fact that the EU Climate Law does not include provisions about the role of international credits post-2040.
17. With a leverage ratio of 1:5, the coalition would offset $15 \times 0.45 \text{ Gt} = 6.75 \text{ Gt}$ and stimulate $6.75 \text{ Gt} \times 5 = 33.75 \text{ Gt}$ in partner countries.
18. The net global emission reduction would be $33.75 \text{ Gt} - 6.75 \text{ Gt} = 27.0 \text{ Gt}$.
19. The estimate is based on an extrapolation that assumes that the current decennial use rate of the remaining 1,150 Gt of carbon budget (2020-2030: 37 percent) will linearly decrease to 0 percent from after 2070 (Zachmann 2025).
20. Collateral requirement = number of borrowed units x leverage rate/expected mitigation value = $100 \times 5/0.5 = 1,000$.
21. See Swiss Federal Office for the Environment, [‘Bilateral climate agreements’](#), 21 November 2025; see Ministry of Foreign Affairs of Japan, [‘Joint Crediting Mechanism \(JCM\)’](#), 9 April 2026.
22. The division between buyer and seller countries does not have to be static. As seller countries develop, their abatement costs may rise, making sales of carbon credits less attractive. The EU should start the Carbon Buyers’ Club with countries that are potential carbon buyers today and leave open the option for additional countries to join later.
23. The Article 6.4 Supervisory Body and the Methodological Expert Panel have developed a variety of standards addressing some these issues, but international carbon crediting remains fragmented (Fuchs et al 2026).
24. Such as i) per-capita distribution of future budget; ii) per-capita distribution of 1900-2100 budget minus historic emissions; iii) inverse to GDP per capita. On the impact for the EU of different allocation rules, see ESABCC (2023).

25. One way to approach this trade-off could be multiple levels of baselines. Every tonne of carbon mitigated compared to a tight baseline would be accounted for as a full tonne of mitigation, while staying below a looser baseline would only be counted as a fraction of a tonne of mitigation. Again, properly calibrating the set of baselines to maximise incentives and minimise windfall profits requires a defensible methodology.

References

- Asadnabizadeh, M and E Moe (2024) *'A review of global carbon markets from Kyoto to Paris and beyond: the persistent failure of implementation'*, *Frontiers in Environmental Science* 12: 1368105.
- Bahí, A, M Fuchs and C Reverdy (2026) *'Carbon pricing beyond borders: assessing climate policy spillovers from the EU carbon border adjustment mechanism'*, Working Paper 05/2026, Bruegel.
- Cames, M, R Harthan, J Füssler, M Lazarus, C Lee, P Erickson and R Spalding-Fecher (2016) *How additional is the Clean Development Mechanism?* Oeko-Institut.
- Crippa, M, D Guizzardi, F Pagani, M Banja, M Muntean, E Schaaf ... F Pekar (2025) *GHG emissions of all world countries*, Joint Research Centre, European Commission.
- CRS (2025) *U.S. Withdrawal from the Paris Agreement: Process and Potential Effects*, CRS Report R48504, U.S. Congressional Research Service.
- Delbeke, J and J Bencini (2026) *'The state of carbon credit policy post-COP30'*, Policy Brief 2026/01, Florence School of Transnational Governance.
- Edenhofer, O, C Leisinger, L Stern and M Kalkuhl (2026) *'Making international carbon markets work for Europe: Jurisdictional Reward Funds and the EU's 2040 climate target'*, PIK Policy Paper, 16 June, Potsdam Institute for Climate Impact Research.
- ESABCC (2023) *Scientific advice for the determination of an EU-wide 2040 climate target and a greenhouse gas budget for 2030–2050*, European Scientific Advisory Board on Climate Change.

- European Commission (2026) *'Proposal amending Directive 2003/87/EC and Decision (EU) 2015/1814 as regards driving competitiveness and cost-effective decarbonisation'*, COM(2026) 616 final.
- Fuchs, M, T Mramor and G Zachmann (2026) *'Buyer beware: a taxonomy of the risks of international carbon credits'*, Analysis 17/2026, Bruegel.
- Glennerster, R and S Jayachandran (2023) *'Think Globally, Act Globally: Opportunities to Mitigate Greenhouse Gas Emissions in Low- and Middle-Income Countries'*, NBER Working Paper 31421, National Bureau of Economic Research.
- McWilliams, B and G Zachmann (2021) *'Commercialisation contracts: European support for low-carbon technology deployment'*, Policy Contribution 15/2021, Bruegel.
- Montrone, L, J Leukers, S Jha, S Scholz, A Wehrl and H McDonald (2026) *'Article 6 credits in the EU ETS 1: Brief assessment of possible integration options'*, Discussion Paper 44/2026, German Environment Agency.
- Nicolini, M and M Tavoni (2017) *'Are renewable energy subsidies effective? Evidence from Europe'*, Renewable and Sustainable Energy Reviews 74: 412-423.
- Pisani-Ferry, J, S Tagliapietra and G Zachmann (2023) *'A new governance framework to safeguard the European Green Deal'*, Policy Brief 18/2023, Bruegel.
- UNEP (2025) *Off Target, Emissions Gap Report 2025*, UN Environment Programme.
- Zachmann, G (2025) *'A Strawman Proposal to Use International Flexibility in Achieving Developed Countries Climate Targets to Catalyse Global Decarbonisation'*, The Economists' Voice 22(2): 339-348.

The authors thank Maximilian Fuchs, Flora Marchioro, Alexander Roth, Simone Tagliapietra and the participants in the Bruegel roundtable on international carbon trading on 18 June 2026 for their helpful comments. This article is based on *Bruegel Policy Brief* Issue n°16/26 | April 2026.



Download EQ Europe Quarterly Autumn 2026, a fully interactive review of European trade and finance for the last quarter.

Download



**SCAN TO
DOWNLOAD**

Securing European prosperity

In the State of the Union address, Ursula von der Leyen shared the key priorities and flagship projects for the year to come, building on the progress the EU has made in recent years and setting out the priorities for the year ahead and her vision for a stronger, more independent and resilient European Union

It was the best of times, it was the worst of times.' This is how Charles Dickens describes Europe in the days of the French Revolution. And it is also a fair description of the world we live in. The speed of change is incredible. The challenges exceptional. So if I look at Europe today, my take is this: the state of our Union is the strongest it has ever been. The state of our Union can also feel as precarious as it has ever been.

These two statements may seem incompatible. But they are both true. And they reflect these exceptional times. This complex age of great possibilities and great perils, which all countries are seeking to navigate. In this world, Europe has chosen to forge its own path. And a stronger, independent Europe is now emerging.

This is a Europe that is securing its prosperity. One that is breaking free from its toxic dependency on Russian fossil fuels. That has become the greatest trading power on the planet. That is transforming its industrial policy while upholding our unique social model.

This is a Europe better able to defend itself than ever before. One that has radically overhauled its approach to securing our homelands. That has increased defence spending by almost 80 % in the last five years alone.

This is a Europe that stands up for each other. As we are doing by standing with Ukraine as it fights for its future and for our freedom. Or when we stood up for Greenland when it was threatened. Or when drones, hybrid attacks or disinformation target our democracies.

All of this proves one thing. That Europe is there for Europeans. And that in the cold fragility of today's world, only Europe can provide true sovereignty to Europeans. This is a Europe we can be so proud of.

At the same time, we all know the threats that we are facing. As old assumptions and relationships have frayed, we are now dealing with an openly hostile world. A territorial war of aggression is raging on our continent; others have

erupted nearby. A fight for industrial capacities is shaping global competition. And a battle of narratives is seeking to weaken confidence in Europe; to paralyse our democracies.

But people's worries are not only linked to these big, global threats. The effects of the many crises are present for many European families, not least on the cost of living or housing. The dizzying speed of change and technology is impacting our jobs, our livelihoods and our society. And it is also having a profound effect on our democracy and our discourse. We see this in the rise of the extreme, the decline of nuance – even in the way that we relate to each other.

*In this world, Europe has chosen to forge its own path.
And a stronger, independent Europe is now emerging*

These real concerns are being weaponised to drive a wedge between us. Sometimes from the outside – and too often from the inside. But our path is clear. Building an independent Europe that has the power to act. Nothing can divert us from this. Because in this world, the fate of Europe must remain in the hands of Europeans.

So our choice is clear. Do we want a strong and independent Europe that stands up for its values in this world? Or do we let divisions and dependencies hold us back? Do we want to rally around our European idea that together we can decide our own destiny? Or do we let our democracies be undermined by the proxies and puppets of authoritarians?

Whether it is the ultranationalists or the anti-Europeans, whether it is Russian interference or disinformation, the names may be different, but their goal is the same. They despise our values, and they want to shatter our European unity. So let's not only fight back against that, let's fight together for our European idea.

Our economy has withstood the impact of the Middle East wars better than expected. And unemployment is at near-record lows. But the pressures from higher energy prices and borrowing costs are biting for people and for businesses. Clear fiscal choices will have to be made to ensure fiscal sustainability, while protecting investment. So, kickstarting Europe's economy is our number one priority.

We have a huge market, world-class industry and services, a rich pool of savings and one of the world's leading currencies. And yet, we still fragment our capital, networks and purchasing power. This is why, when we took office, we put in place a bold plan. To build an economy where innovation and energy flow across borders. Where companies compete and grow across Europe, and savings finance our future. This is the logic of the Draghi and Letta reports. And together we are delivering. But we have to be faster – all institutions.

We now have the political agreement on the One Europe, One Market Roadmap. And together, by the end of next year, we can and must complete this historic overhaul of the Single Market. In today's economy, speed is of the essence. A permit or a form, a grid connection or a financing decision. All of this can determine where an investment is made. Where jobs are created. We need to move much faster. This is why we will put forward proposals to further speed up permitting.

Where projects are in Europe's strategic interests, we need to massively accelerate. And where the administrative burden is slowing down business and SMEs, we need to slash that too. You know that our 12 omnibus proposals are reducing the administrative burden by around €17 billion a year. But simplification cannot be a one-way street.

All levels have to have skin in the game. We have long spoken about the harms of gold-plating. Yes, we are doing our best at the European level to cut red tape. But now I think it is time for the member states to step up. If we are serious about simplification, we also need a pact against gold-plating.

Europe cannot remain an industrial powerhouse if its energy prices are structurally too high. When Russia cut off the gas we fought back. Our supplies are now more diversified. And we have invested in our homegrown clean energies – renewables and nuclear. But then came the closure of the Strait of Hormuz. And again we bitterly felt the cost of our general dependence on imported fossil fuels. Since the start of the conflict, imported fossil fuels have cost us an additional €90 billion, without a single molecule of energy added.

So, we need to double down on our affordable, homegrown, clean energy. Be it renewables and nuclear, or biomethane and others. Technologically neutral. But they must give us independence. For this independence, we need to electrify Europe. Our goal is to double the share of electricity by 2040. With this, Europe could cut its fossil-fuel import bill by €260 billion a year. But we also need to ensure that the electricity we generate is available.

Last year, we installed more than 80 gigawatts renewable capacity. But six times more capacity is still waiting to be connected. This is why we need the grids package so urgently. We must invest faster. Speed up grid connections. Develop storage. In short, let's make Europe's future electric.

We can bring down energy prices, invest and innovate more, reform our economies. But all this is undermined if our companies do not compete on a level playing field. Our trade deficit with China is now €1 billion – a day. It has reached a tipping point. Some say the second China shock is looming. But it's already here. It shows in our communities and in factories across our Union. It leads to deindustrialisation in the industrial heartlands of Europe. This is unsustainable.

And this is why we are engaged in a dialogue with China to rebalance our trade. But this dialogue must now lead to results. China's weaker domestic demand means it also needs our European market. So, it is in both our interests to work together to find solutions. Let me be clear: we will use all the tools at our disposal to rebalance our relationship. Words are good. But deeds are better.

And there is another point. We still have too many dangerous dependencies. We are more than 80 % dependent on China for many critical raw materials. 90 % for some rare earths. We have done a lot. But we are not the only ones trying to diversify. We need to urgently procure and build up our reserves. No country can do this alone. So we need to think differently.

This is why we will establish a new European Corporation on Critical Raw Materials. It will help us obtain and stockpile what we need. For electric cars, chips and batteries, clean tech and defence. And so much more. Not only our industry depends on this, but also our independence and security.

This brings me to my next point, on financing. To be very clear, our budget is about our future. And any modern budget needs new own resources. We must create the conditions for our companies to attract investment and grow. This is why our work on the Savings and Investments Union is so critical. But developing capital markets takes time. And the needs of our companies are urgent. We need a banking system that is built for growth – not just stability.

This is why we will put forward a new Banking Package focusing on simplification and tackling fragmentation. We know the demand for capital is out there. But we need a greater capacity and appetite for risk. Our deep-tech start-ups are on track to attract record venture capital investment in 2026. This is why I announced a new Scaleup Europe Fund last year.

A year on, the first investments have been made in companies such as ICEYE. ICEYE was set up by two young Europeans: Rafał from Poland and Pekka from Finland. They initially met through Erasmus some 15 years ago. Together they founded a tech start-up. They developed their idea around satellite imagery. They got Horizon funding. They proved their technology. They grew with private investment.

And today they are one of the world's leading space technology companies. With offices in Finland, Poland, Spain and Greece. ICEYE's story shows that Europe is the home of innovation. And that we can reignite Europe's entrepreneurial spirit.

These are Europe's economic fundamentals. But our future prosperity and security will greatly depend on how we manage the tipping points of our time: climate change and AI. On climate change, this summer was the summer of truth. Almost no part of our continent was spared. In Belgium's High Fens, a peatland formed over centuries burned into a wasteland of ash. Flames roared from the mountains of Ireland to the islands of Greece and Croatia. In the

Alps, ancient glaciers retreat further every day. The impact of all of this is simply devastating: 660 000 hectares burnt to the ground. Some 12 million tonnes of crops lost.

Key rivers like the Danube, der Rhein and La Garonne are disappearing before our eyes. More than 35,000 additional deaths during the heatwaves. Care homes and hospitals dangerously overheated. And all of this is a mere preview. These were the hottest summer months ever – but probably the coolest of the years to come. Science is clear. Europe is warming at double the global rate. Of course, the transition is complex. And we will need to adjust and learn along the way. But there is no doubt: Europe can, must and will stay the course on its climate targets.

This is not about mitigation or adaptation. We need both. And they need to strengthen each other. This is the only way to reach climate resilience. But we need to step up dramatically on preparedness. We will present a new framework for climate resilience next month. We will identify 100 of the most vulnerable territories in Europe. And we will assess the risks – and at which level they should be managed.

This will give us a whole of society approach – all the way through to the local level where adaptation matters most. And it will aim to make Europe the champion in adaptation technologies – a huge emerging market. Whether drought-resistant crops, flood prevention or energy-efficient cooling, Europe is already a leader. And we must capture this market.

Today, only around 25 % of catastrophe losses in Europe are covered by private insurance. This means that, far too often, national budgets become the insurer of last resort. We need to take action to close this gap. This is why the Commission will set up a Climate Insurance Alliance. There can be no repeat of this summer. This is why we will also soon present a European Heatwave Plan. We need better early warning systems and better health preparedness. We need urban adaptation and cooling. And we need to support the most vulnerable.

The second point is on water and food security. We know the risks of water scarcity for our farmers, energy and supply chains. Yet, European pipes lose almost one in every four litres Mostly through leakage. What a waste. So we will put forward a new European Water initiative. Because water is precious. For our industry – and for our food security. Crops need water to grow, while livestock need it to survive. Entire harvests depend on it.

This is why we will make proposals to improve risk management and address the insurance gap. Once again, Europe's farmers have risen to the challenge of this never-ending summer. Europe's farmers are the best in the world. And we will support them every step of the way.

The next point is about wildfires. They are getting more intense, more frequent and more spread out. I am so proud of Europe's response. Of our European assets flying across the continent. But clearly, we need to think bigger. And learn faster. This is why we have asked some of Europe's most senior firefighters and emergency responders to lead the work.

They will look at all aspects of preparedness and prevention. From civil–military coordination to needs for new capabilities. And I have asked them to make proposals, including for a future European Firefighting Fleet. Our emergency crews have saved countless lives and homes this year alone.

The second tipping point of our times is artificial intelligence. AI is fast becoming a foundational layer for our economy and security. Europe's stance is clear. We want to make the most of AI to improve our society, our quality of life and our productivity. I am an optimist that AI has the potential to benefit humanity – if we get it right. But, as with any powerful new technology, AI comes with a balance of possibilities and risks. Without addressing these risks, we will never be able to unlock AI's possibilities.

As the frontier models become more capable, these risks have come more sharply into focus. Models being developed will allow hacking on a level we never thought possible. And they will soon be in the hands of adversaries who see the world very differently from us. At the same time, the dangers of self-improving models are becoming ever more apparent.

Incidents of AI agents escaping their environment or inserting malicious code are a mere glimpse. After the HuggingFace incident, developers are ringing the alarm. And CEOs of the most advanced companies tell us that it is time to slow down on the self-recursive models. To pace the frontier.

If the people developing the technology are clear, then we should be too. The AI Act is of course a crucial piece to putting guardrails in place. And it puts Europe in the position to shape global efforts on how to tackle this. So, we will step up together with our like-minded partners like Canada, the United Kingdom and others. We want to team up on model evaluation, verification, early warning, AI security and much more. And I will invite the main frontier labs for a discussion on how we can support ongoing industry efforts to pace the frontier.

It is also true that Europe will need its own capabilities – in particular to protect our national security. We need to stay in the race. Because this is crucial for our independence. So we need to massively boost our computing capacity. And we will set out our ideas for how to achieve this in a clean and efficient way. We will also develop new ways to use public and private funds to invest in the most promising European start-ups and companies.

By limiting the risks of AI, we can maximise the benefits. We have already done a lot. On gigafactories and tech sovereignty, for example. But the next leg of this race is not only about frontier models. Just as was the case with the printing press in times gone by. It was invented by a German inventor, but was mastered and applied across Western Europe.

We do not need to be the ones who develop the frontier technology to be the ones who draw the greatest value from it. The key is creating the value that AI has not yet delivered. Moving AI from the screen to the real economy and society. On the factory floor. On the hospital ward. In precision agriculture. In our energy grids. For autonomous driving. For defence. This is where the real value is hidden. And it is where Europe has big advantages. We have world-class industries that sit on the highest-quality data. These data are a precious European treasure. They are what can power tailor-made industrial AI models.

Let me give you one societal example. On breast cancer screening – an issue many of you have campaigned tirelessly on. Mammography screening reduces mortality by around 30–40 % among women who attend the screenings. Early detection is of the essence. This is where AI can step in. AI-supported mammography can identify earlier and more cancers during screening. And as a result more women can survive. So we need AI for health. But, here again, we must do it in a European way.

Do I want my doctor to have instant AI access to all the data needed to make the best diagnostic or treatment choice? Yes, of course. But do I want my doctor to be a robot, AI assisted? No, of course not. I do not want a robot to tell me whether I have cancer or not. The point is that AI needs to empower our doctors. Not to replace them.

This is just one example. We are now focusing on five of the highest-value sectors in which industrial AI has the biggest potential and we have the data: health, transport, agri-food, advanced manufacturing, and defence and space. In November, we will announce game-changing initiatives across these sectors. We will invite member states, MEPs, industry and entrepreneurs to join us in this ambitious European project. Our approach is clear. We want AI. We want AI to create more value more responsibly. And at a lower cost and with less energy consumption. But above all, we want AI with human agency.

This is our European way. Many of the risks of AI also spill over into our social market economy. For some, AI allows them to speed up repetitive tasks, to make better use of their skills. But clearly, these benefits are unevenly spread between skill sets, generations and regions. Look at young people and the barriers they are facing, just to get on the first rung of employment.

Our response must keep pace. From education to employment, AI must help develop skills, improve the quality of jobs and give workers a fair chance. This is why we need our Quality Jobs Act. And we will deliver before the end of the year.

Next year, we will celebrate the 10th anniversary of our European Pillar of Social Rights. And we will keep delivering on its promise. We will invest in our regions, so everyone has the 'right to stay'. We will launch a European Care Deal to address the challenges linked to our demographic changes. We have just proposed the first-ever Housing Act to tackle short-term rentals and housing speculation.

Staying in the region is not a privilege. Care is not a privilege. Housing is not a privilege. They are rights. When we proclaimed this Pillar of Social Rights 10 years ago, the world was different. Our society has changed. In many ways for the better, in others not. Today, our fundamental values are being challenged. Some people even want to backslide on women's rights. To send us back to an age in which women were second-class citizens.

We have fought so hard for equal treatment, equal opportunities, equal pay. So let me send a very clear message to all those who need to hear it: women are here to stay. We will defend our rights. And we will continue to stand up for our values every single day. This is why, together with my dear friend António Costa, we will host a summit to uphold the enduring principles of our Social Pillar. Times may change, technologies may change, but our values do not. In Europe, people will always come first.

In this age of tipping points, our prosperity and security will not only be built at home. This year, we have signed free trade agreements with India, Mercosur, Mexico, Australia and Indonesia. We now have trade deals in place with more than 80 countries – more than anyone else in the world. And business is asking us for more.

I can announce a new international connectivity project. It will link the South Caucasus and Central Asia directly to the European market – what we call the Middle Corridor. Through Global Gateway, we will aim to crowd in up to €12 billion in public and private investment. Our goal is to diversify routes, triple trade flows and slash freight transit times by 2030. To make this a reality, we will organise a Regional Connectivity Summit with the Prime Minister of Bulgaria.

These partnerships are a strategic choice for Europe. But they also respond to the fracture in the international rules-based system. For some, the answer is to go it alone. But this is not an option for us. Europe will always lead when it comes to defending the rules-based system. But we can no longer rely on it as the only way to secure our interests. Or assume that its rules will shelter us from the complex threats that we face.

In this new world, we must urgently reimagine our partnerships. And build global coalitions for our resilience and democracies. And here too only Europe – with its size and power – can take the lead. Thanks to CETA, our trade in goods with Canada has grown by 75 % in less than a decade. This shows the power of our trade agreements.

We see the world with the same eyes: from AI to climate change, from the Arctic to geopolitics. And we have stood together: on Ukraine, on defence, on raw materials and on supply chains. We believe that power does not belong to the strongest, the richest or the loudest – but to all of us. That democracies have the freedom to choose with whom to work.

So, we are joining forces voluntarily. We will move from CETA to an Alliance for the Future to create a common prosperity and economic security space. We will work on intelligent manufacturing. We will create a tech alliance. We will integrate defence industrial bases. We will make the Arctic a flagship joint project. We will work on energy, critical minerals and batteries. On AI, quantum, cyber and economic security. This is a partnership not against anyone else, but for our common strength. In short, we want to bring the relationship with Canada to the highest level possible.

I said we must urgently reimagine our partnerships. So, I would like to work with you on opening the door for Canada to be the first associate member of the EU. We share one ocean, one set of values, one way of seeing the world. And we will now build our shared future as well – thank you for being here.

Europe's belief in cooperation has always been driven by peace. It is our *raison d'être*. Today, that peace is threatened. Russia's war of aggression on Ukraine rages on. And threats are mounting on our soil. In the last week we have seen it in Denmark, Lithuania and Poland. The attempted drone attack in Leipzig is further proof. We need to be clear about what this was. An attack using military-grade material, carried out by Russian operatives on European soil. An attack not only on one member state – but against all of our Union. And, we will not allow it.

We should be under no illusions about what is happening here: hybrid threats facing Europe are less and less hybrid and less and less threats. Cyberattacks and sabotage, arson and drone incursions. They are ramping up every day. So, as the threat level climbs, so must Europe's preparedness. This is why I will propose, together with the HRVP, a new European security strategy.

We already have many tools. And we are strengthening our work with NATO. But our doctrine, our institutions and our decision-making have not kept pace with the new reality. Our systems were made for a different world. They are

driven by well-intentioned attempts at consensus and compromise. But we need to consider whether they are still fit for purpose.

For instance, Europe needs its own counter-hybrid playbook. We urgently need a consensus on how to respond to certain incidents. Those which fall below armed aggression but clearly threaten our national security. In other words, a mechanism to flank NATO's 'Article 4'. I believe it is time for Europe's own Article 4 – or an Emergency Security Protocol. When triggered by one member state, all member states would convene. To coordinate a European response. To deter further escalation. And to mitigate the impact.

And I believe we must go further still. For a long time, we have discussed a leaders' format focused on the security of our continent. With partners like Canada, Norway, Ukraine, the United Kingdom and others involved. This is even more urgent now, given the nature and scale of the risks at play. This is why we will set out our ideas to make a European Security Council a reality. The urgency is clear.

This united front also drives our work on European defence. The surge in defence spending has led to record levels of investment in our capabilities and forces. New joint projects are getting off the ground. And we are spending more European than ever before. Bigger orders for our industry, more interoperability and more economies of scale. Now we need to take that approach to the next level.

The war in Ukraine has shown us the importance of strategic enablers. These are the critical support assets and force multipliers that any credible defence relies on. From air and missile defence to strategic transport and air-to-air refuelling. From space to cyber. These systems are essential, and we need more of them. Only together does Europe have the scale to deliver them. That is why I can announce a new European Instrument for Strategic Enablers – fully aligned with NATO. Because only a strong and credible European defence posture will guarantee our security.

And of this there is no room for doubt: Europe will defend each and every single square metre of its territory. Every day this summer, strike after strike hit Ukrainian cities. In Kryvyi Rih, Russian drones struck a crowded shopping mall in broad daylight. Then they struck a second time – deliberately targeting the rescue workers who rushed in to help. This is not just war. This is a deliberate strategy of terror. But Ukrainians have proved countless times that their spirit cannot be broken.

The deep suffering in Gaza and unchecked settler violence in the West Bank have deeply affected Europeans. The Israeli government's decision to advance the illegal E1 settlement project is unacceptable. Last year, I called for Europe to act, proposing a number of measures. It is painful for many Europeans that member states have been unable to build the necessary majorities for a unified response. Some member states have moved ahead. Others are considering doing so.

When Europe is divided, it cannot change the course of events. The months ahead in the region will be critical. Despite these challenges, Europe is leading the way in supporting Palestinians. We have delivered more aid than anyone else. Through the Palestine Donor Group and the Team Gaza Initiative, we have brought together 65 delegations. Together we have pledged nearly €1 billion to rebuild Gaza and support a modern Palestinian Authority. Europe supports the right to self-defence for Israel. And the Palestinian people have an equal right to live in safety and with dignity. It is Europe's role to keep the flame of a two-state solution alive.

I have spoken about the need for Europe to act together. Or risk being pulled apart. Once again, this summer, this reality was tested at our borders. From our eastern borders to the shores of Crete and right across the Mediterranean. And of course, with the crisis in Ceuta. These events are different in nature. But they have two things in common. The first is that vulnerable people are being exploited. By callous traffickers and smugglers. By our adversaries. This is why we have proposed a sanction regime to hit criminals where it hurts most.

This brings me to the second parallel. All of our external borders are European borders. So, I want to address the people of Spain directly. Ceuta's border is a European border. Because Ceuta is Spain. Ceuta is Europe. And Europe will be there for you. The point is that the only solution is a European solution. And that is the Pact on Migration and Asylum.

After years of work, we finally have a common European framework. One that gives responsibilities to all – and that demands solidarity from all. And it now needs to be fully implemented. We are on the right path. We have seen a 55 % reduction in illegal arrivals over the past two years. And a further 35 % drop this year. That is real progress.

But the events of this summer show that our tools need to evolve – especially for emergency situations. Europe needs the means to protect its borders at all times. Especially in scenarios when mass movements are engineered and weaponised. This is why we will propose a new European Emergency Response Framework. It will only be triggered under a strictly defined set of criteria. And in those cases, it will allow for time-limited and strictly defined flexibility in standard procedures.

The message from Europe is clear: we will always comply with our values and fundamental rights. But we will never compromise on protecting our borders. It is we, Europeans, who decide who comes to our Union and under what circumstances. And not the smugglers and the traffickers. This is why we will also step up our work to strengthen Frontex – especially to speed up returns. We will improve our early warning systems – including in partnership with third countries. We will ensure that we can detect and anticipate online. And finally, we will keep tackling the root causes of illegal migration.

We need to work together with our neighbours to boost investment. To offer young people meaningful prospects – so they are not compelled to gamble with their lives. Jobs, skills, training, apprenticeships. This is why I am announcing a new Mediterranean Youth Skills and Jobs initiative.

This is how we can reduce illegal migration. Uphold the rights of individuals seeking protection. Facilitate safe, legal migration with new Gateway Offices in partner countries. And defend the freedoms of our Schengen area. But above all, this is how we prove that Europe can act together.

This need for collective European solutions is more important than ever. Because it is also our European democracy that is being threatened. Time and again, we have seen disinformation and foreign interference seeking to divide us. Seeking to erode trust in our democratic fabric and free media. Election campaigns poisoned by foreign propaganda. Conspiracy theories built on viral deepfake videos.

This is cognitive warfare. In Europe, we form our own opinions based on access to reliable, independent information. This is how we debate, compromise and find solutions – even on controversial issues. This openness is our biggest strength. But it is being weaponised by our opponents to distort debate and truth.

So we need a collective capacity to anticipate, detect and respond. This is why, last year, I announced the Centre for Democratic Resilience. It is already delivering with all the member states. We will step up our work, pooling our resources with member states and stakeholders. But we should be under no illusions. The threats to our democracies do not only come from the outside. We see the rise of anti-democratic, anti-European forces. And we know the risks.

But nothing is inevitable. Just look at Hungary. The date of 12 April will stay in our memory for a very long time. The day the Hungarian people took their future into their hands. They chose democracy. They chose Europe. They brought Hungary back to where it belongs: at the heart of our Union. Years of backsliding have started to be reversed. The fight against corruption and state capture is speeding up. There is real progress on fundamental rights. On academic freedom. And now, billions of euro of investment in Hungary have been unlocked.

There is of course much more work ahead. But this shows that hard work pays off. Thanks to our Rule of Law Reports, any issues are now detected early and addressed through dialogue. And we have strengthened the link between funds and respect for the rule of law. With the next long-term budget we must go even further. Respecting the rule of law and fundamental rights is a must for EU funds. It has to be front and centre in the budget. Now and in the future.

At its heart, democracy is about giving people the power to choose the society they want to live in. Today, much of this power has been taken out of the hands of parents. Young Europeans now spend four to six hours per day on screens. Children are pulled ever deeper into feeds designed to keep them scrolling. What we are witnessing is a great capture of our children. A capture of their attention, their focus, their minds. This is depriving children of their childhood.

What our children need is time to grow into themselves. To grow up with others. To be confronted with all of the emotions and choices that come with real life. They even need time to get bored. Because this is when they have to use their imagination. When they have to talk, agree, disagree, form an opinion, change an opinion. But when a child has a smartphone, all of this is taken away.

Day and night, parents see the cost: loss of sleep, anxiety, even self-harm. And in a growing number of cases even fatal tragedies. C'est trop. Enough is enough. I convened a Special Panel of experts. We spoke to young people. We looked at Australia and others who have moved ahead. And now it is time for Europe to act. The Commission is proposing the EU KIDS ACT. We will take a gradual and differentiated approach. Each age has its own sensitivities. And therefore, each will have its own tailored protection level. No social media under the age of 13. No personal account under the age of 15.

That means, from 13 until a child turns 15, only mini accounts set up and supervised by parents, with limited features and time restriction. And between 15 and 18, safe design will be an obligation for the platforms. I am aware that many perceive the power of big tech as overwhelming. And impossible to roll back.

I disagree. We do not have to accept addictive features. We do not have to accept children being drawn into ever more extreme content. We do not have to accept that girls have their photos used for AI-generated sexualised images. So we are reversing the burden of proof.

Platforms will have to prove that they are safe. Because it is not about our minors accessing social media. It is about when and how do we allow social media to access minors. Europe has the power to act. It is we who decide our rules, not big tech. But it is not only minors who are at risk. Addictive design, for example, is harming everyone. This is why we need a wider framework, the Digital Fairness Act.

Long live Europe. ■

Ursula von der Leyen is President of the European Commission

This article is based on the [2026 State of the Union Address](#), Strasbourg, 16 September 2026.

Rebuilding the world order

As the postwar international order fragments and universal agreement becomes increasingly elusive, policymakers face a pressing question: what form of cooperation can replace the postwar hegemon-led, rules-based system? Dennis Snower examines two complementary answers

The postwar international order is no longer functioning as its architects intended. Rules remain, but their authority has weakened. Institutions remain, but their capacity to organise collective action has declined. Economic interdependence remains, but it is increasingly treated as a source of vulnerability rather than only mutual gain.

Recent VoxEU contributions describe this condition as geoeconomic fragmentation: a world in which trade, finance, investment, supply chains, technology, and payment systems are increasingly shaped by geopolitical rivalry (Aiyar and Ilyina 2023, Aiyar *et al* 2023, Chari *et al* 2025).

Two recent approaches offer responses to this changing world order. The first is the *middle-power strategy* of Mark Carney, articulated in his 2026 speeches at the World Economic Forum and the Lowy Institute. The second is the *pivotal-power strategy*, summarised in my paper *Toward a New World Order* (Snower 2026), as part of The New Alignment Initiative of the Institute for Global Prosperity at University College London. Both start from the same premise: the old order cannot simply be restored. But they differ in their diagnosis, their theory of agency, and their implications for national policy and international agreements.

The middle-power approach

Carney's argument begins with geopolitical realism. The rules-based international order helped countries such as Canada prosper, but it was always partly dependent on American hegemony. That bargain, he argues, no longer works.

Great powers increasingly use trade, finance, technology, and supply chains as instruments of coercion. Middle powers must therefore stop pretending that the old order still operates as advertised and instead build institutions and agreements that actually function (Carney 2026a).

In his Lowy Institute speech (Carney 2026b), Carney develops this into a strategy of middle-power cooperation based on 'variable geometry': different coalitions for different issues, reducing vulnerability to coercion while renewing international cooperation one issue at a time. The Carney approach is therefore a strategy of middle-power renewal.

The future international order will require both approaches: the middle-power strategy to preserve strategic autonomy and the pivotal-power strategy to organize effective cooperation on transnational challenges

The pivotal-power approach

The pivotal-power approach begins from a different starting point. The central problem is not only the decline of the rules-based order, but the mismatch between global problems and existing governance structures.

Climate change, biodiversity loss, financial instability, food insecurity, cyber insecurity, artificial intelligence, pandemics, and migration are interconnected problems requiring collective action at a scale that universal multilateralism can often no longer deliver (Snower 2026).

The proposed response is the 'New Alignment': a framework for state-enabled polycentric governance. Cooperation is organised through issue-specific 'pivotal power coalitions'. A pivotal power is not defined by its general rank in the international hierarchy, but by whether its capabilities, interests, institutional influence, and legitimacy are indispensable for addressing a particular problem. A country may be pivotal in climate governance but not in digital governance; pivotal in food security but not in finance.

The New Alignment is structured around three forms of alignment. Goal alignment orients cooperation toward human flourishing, operationalised through 'SAGE' (Solidarity, Agency, material Gain, and Environmental sustainability). Capacity alignment organises the capabilities required to address each challenge, including public authority, technology, finance, science, regulatory capacity, and implementation capacity. Policy alignment integrates instruments across ministries, sectors, jurisdictions, and institutions so that policies reinforce rather than undermine one another.

The main differences

The first difference concerns the *unit of agency*. Middle powers (states located between hegemons and small powers) are the key actors in one approach, while pivotal powers (actors whose participation is necessary for

solving a specific problem) are the key actors in the other approach. The former is status-based; the latter is function-based.

The second difference concerns the *organising purpose*. The middle-power approach is principally about strategic autonomy and renewal of the rules-based order. The pivotal-power approach is about organising effective governance under conditions of partial non-cooperation. The first approach examines how middle powers can defend and rebuild an international order. The second approach examines how global problems can be solved when universal agreement is unavailable.

The third difference concerns *coalition design*. Carney's coalitions are primarily among like-minded middle powers. Pivotal powers are problem-specific and may include great powers, middle powers, developing countries, firms, international organisations, cities, investors, and civil society organisations. Coalitions are defined by functional necessity rather than political affinity.

The fourth difference concerns the *normative foundation*. The middle-power framework is grounded in sovereignty, territorial integrity, legality, human rights, and rules-based cooperation. The pivotal-power framework adds a broader evaluative purpose: human flourishing.

This matters because the legitimacy of international cooperation now depends not only on rules, but on whether cooperation improves the lived conditions of citizens: solidarity, agency, material security, and environmental sustainability.

The fifth difference concerns *governance architecture*. The middle-power approach remains largely within a state-centric diplomatic framework. The pivotal-power approach moves toward polycentric governance, in which states

enable cooperation among multiple decision-making centres: governments, firms, cities, universities, development banks, investors, regulators, and civil society. This reflects the dispersion of relevant capabilities across public and private actors.

The sixth difference concerns *enforcement*. The middle-power approach emphasises trust, legitimacy, and coalition-building. The pivotal-power approach adds institutional guardrails: transparency, peer review, shared metrics, conditional access to coalition benefits, dispute resolution, and graduated sanctions.

Implications for national policies

For national governments, the middle-power approach implies a strategy of resilience. Countries should diversify trade, reduce strategic dependencies, build domestic productive capacity, invest in critical minerals and energy security, and cultivate trusted partnerships. Foreign policy becomes credible only when economic vulnerability is reduced. Middle powers must therefore combine domestic renewal with external coalition-building.

The pivotal-power approach implies a broader reorganisation of domestic policymaking. Governments must identify the transnational challenges for which they are pivotal, then align national goals, capacities, and policies accordingly.

Climate policy cannot be left to environment ministries alone; it must involve finance, industry, trade, agriculture, energy, cities, firms, and investors.

Digital governance cannot be confined to technology regulators; it must involve competition policy, education, labour markets, privacy, security, democratic accountability, and social protection.

This implies that national strategies should be designed around problems rather than ministries. Governments need cross-cutting policy architectures, common metrics, investment pipelines, public-private coordination mechanisms, and institutional capacity to participate in multiple coalitions simultaneously. In economic terms, the state's role shifts from correcting isolated market failures to orchestrating capabilities around systemic challenges.

Implications for international agreements

Carney's approach points toward a new generation of middle-power agreements. These would focus on trade diversification, supply-chain resilience, critical minerals, defence cooperation, AI governance, energy security, and trusted investment channels. They would be pragmatic, sectoral, and flexible. This corresponds to Carney's idea of 'variable geometry': not one grand bargain, but many coalitions.

The pivotal-power approach goes further. International agreements should be designed as modular coalition agreements around specific global challenges. They should specify shared goals, identify the capabilities required, align policies across jurisdictions, and include guardrails to maintain compliance. Agreements should remain open to accession, using benefits such as finance, technology transfer, market access, data-sharing, regulatory recognition, or infrastructure participation to attract additional members.

Recent analyses likewise argue that international cooperation must adapt to fragmentation through flexible coalitions and institutional innovation (Aiyar and Ilyina 2023, Campolmi *et al* 2023, Garicano 2025, Hoekman 2024, Chari *et al* 2025).

The implication is that international agreements should become more problem-specific, more capability-based, and more enforceable. They should neither wait for universal consensus nor dissolve into ad hoc transactionalism. The central design challenge is to combine flexibility with credibility.

Conclusion

The middle-power and pivotal-power approaches offer complementary but distinct responses to the same historical moment. The middle-power strategy is a diplomatic and geopolitical response to the decline of hegemony and the weaponisation of interdependence. It tells countries such as Canada and Australia how to act when great powers no longer reliably sustain the rules-based order.

The pivotal-power strategy is a broader governance response to fragmented interdependence. It tells the international community how to organise cooperation when global problems require action but universal agreement is unavailable.

The difference matters. Middle powers can help defend the remnants of the old order and build new coalitions of trust. Pivotal powers can help construct the architecture of the next order: issue-specific, polycentric, capability-based, and oriented toward human flourishing.

The future international order will require both approaches: the middle-power strategy to preserve strategic autonomy and the pivotal-power strategy to organize effective cooperation on transnational challenges. Without the former, states remain vulnerable to domination. Without the latter, coalitions may remain too narrow, too state-centric, and too disconnected from the systemic challenges they are meant to address. ■

Dennis Snower is Founding President of the Global Solutions Initiative; a Visiting Professor at the Institute for Global Prosperity, University College London; a Professorial Research Fellow of INET Oxford, Oxford University; an International Research Fellow of Said Business School, Oxford University; and a Non-Residential Senior Fellow at Brookings

References

- Aiyar, S and A Ilyina (2023), *[“Geo-economic fragmentation: What it means for multilateralism”](#)*, VoxEU.org, 29 March.
- Aiyar, S, P-O Gourinchas, A Presbitero and M Ruta (2023), *[“Goeconomic fragmentation: A new eBook”](#)*, VoxEU.org, 2 October.
- Campolmi, A, H Fadinger, C Forlati, S Stillger and UJ Wagner (2023), *[“A proposal to refocus the EU Carbon Border Adjustment Mechanism on preventing carbon leakage”](#)*, VoxEU.org, 3 March.
- Carney, M (2026a), *[“Principled and pragmatic: Canada’s path”](#)*, address to the World Economic Forum Annual Meeting, 20 January.
- Carney, M (2026b), *[“Prime Minister Carney delivers remarks at the Lowy Institute in Sydney”](#)*, 4 March.
- Chari, A, N Converse, A Mehl, GM Milesi-Ferretti and I Vansteenkiste (2025), *[“Geopolitical tensions and international financial fragmentation: The 28th Geneva Report on the World Economy”](#)*, VoxEU.org, 20 October.
- Garicano, L (2025), *[“Rewiring trade for a warming world: How COP30 can turn climate goals into trade rules”](#)*, VoxEU.org, 15 November.
- Hoekman, B (2024), *[“Non-economic objectives, globalisation, and multilateral trade cooperation”](#)*, VoxEU.org, 11 September.
- Snowder, DJ (2026), *[“Toward a New World Order: Reconceiving The New Alignment”](#)*, manuscript.

This article was originally published on [VoxEU.org](#).

The background of the slide is a dark, blue-toned image. It features a world map in the background, with the continents visible in a lighter shade. In the foreground, two hands are shaking, symbolizing a deal or agreement. The hands are wearing dark suits with light-colored cuffs. The overall tone is professional and global.

Redrawing the global trade map

Liberation Day raised an immediate question in policy circles: would other countries follow suit? Barthélémy Bonadio, Andrei Levchenko and Nitya Pandalai-Nayar study the implications of the US withdrawal from free trade

The sweeping tariffs announced by the US on 2 April 2025 – labelled ‘Liberation Day’ by the administration – raised an immediate question in policy circles: would other countries follow suit? Two competing views emerged. One held that countries would be compelled to raise their own barriers to shield domestic industries from trade diversion, particularly from Chinese exports redirected away from the US market. The other held that the mutual gains from free trade are sufficiently large that the rest of the world would sustain the cooperative equilibrium even without its largest participant.

The debate can be framed within a well-established framework in international trade theory. A well-known result is that from any individual country’s perspective, the welfare-maximising import tariff is positive: it improves the terms of trade at the expense of trading partners (Johnson 1951). Quantitative assessments place unilaterally optimal tariffs at 20–62% (Ossa 2014, Costinot and Rodríguez-Clare 2014).

Prior to the first Trump administration, observed tariffs were far lower, typically 3–4% for most country pairs (Bagwell *et al* 2016). This gap is rationalised as the outcome of a repeated prisoner’s dilemma game: the long-run loss from a trade war when partner countries retaliate deters unilateral defection, sustaining free trade as a cooperative equilibrium¹. Using a quantitative model, our paper (Bonadio *et al* 2026) examines what happens to that equilibrium when the US steps outside it.

The analytical framework

To organise ideas, we first work with a basic multi-country international trade model that admits a closed-form solution for unilaterally optimal tariffs. The optimal tariff for any country is a function of a single sufficient statistic: the export-weighted average of foreign countries’ expenditure shares on the country’s goods.

Intuitively, a country's incentive to impose tariffs depends on how much market power it has in the goods it exports — which is measured by how concentrated other countries' import spending is on its output. For most countries, this share is small, and optimal tariffs are close to the inverse of the trade elasticity regardless of the configuration of trading partners. We extend prior analytical results (Gros 1987, Lashkaripour 2021, Ignatenko *et al* 2025) by incorporating endogenous foreign labour supply, which moderates the terms-of-trade gains from tariffs and lowers optimal tariffs for large economies.

In the period since Liberation Day, global trade cooperation does not appear to be decreasing. Several significant trade agreements have been announced or concluded since April 2025: the EU-Mercosur agreement, the EU-India trade agreement, and the expansion of the Comprehensive and Progressive Agreement for Trans-Pacific Partnership (CPTPP)

The second component of our analysis is a quantitative multi-country, multi-sector model featuring a global input-output network and trade (Bonadio *et al* 2021, Huo *et al* 2025), calibrated to OECD inter-country input-output data for 42 countries including an aggregate EU. We use this framework to compute unilaterally optimal tariffs and Nash equilibrium tariffs — which are all countries' best responses to each other — both before and after a US exit from the free-trade regime.

The sustainability of cooperation is measured by the minimum discount factor required to support the cooperative free-trade equilibrium under a grim-trigger strategy as in Mei (2020). A higher minimum discount factor indicates that the cooperative equilibrium is harder to sustain.

Effects on optimal tariffs and cooperation sustainability

Our central finding is that even a complete US exit from global trade cooperation has limited quantitative impact on other countries' optimal tariffs and on the sustainability of trade cooperation among them.

The mechanism is transparent from the analytical results. The US accounts for approximately 14% of world imports, so its withdrawal rearranges bilateral trade patterns worldwide. However, the key sufficient statistic — the export-weighted foreign absorption share — changes little for most countries. The redirection of trade flows is spread across many trading partners, leaving each country's exposure to any individual partner, and thus its incentive to impose tariffs, largely unchanged.

Quantitatively, in the baseline single-sector model, unilaterally optimal tariffs for China and the EU rise by less than half a percentage point following full US withdrawal. The minimum discount factor required to sustain cooperation rises from 0.425 to 0.429 for China and from 0.433 to 0.438 for the EU. For the broader sample of other countries, changes are correspondingly smaller.

These results are robust to changing the trade elasticity and the labour supply elasticity. A multi-sector model with global production network linkages, which features sector-specific optimal tariffs across 16 tradeable sectors and 41 trading partners, implies lower average optimal tariffs due to the second-best effects of input trade, and a somewhat higher median minimum discount factor of 0.476. However, the qualitative conclusion remains the same in this richer environment: the change in cooperation sustainability from the US exit is modest.

Scenarios with partial US exit


We also consider a more realistic scenario in which the US imposes the import-weighted average Liberation Day tariff of 21% rather than withdrawing entirely. The effects on cooperation sustainability outside the US remain small: the average minimum discount factor rises by approximately 0.01 relative to the 2018 baseline.

Figure 1 traces the minimum discount factor for China and the EU as the US tariff varies continuously from zero to 200% and reveals an inverted-U pattern. As US tariffs rise from zero, the required discount factor increases because the US remains a participant in the tariff game and countries face US retaliation if they deviate.

However, higher initial US tariffs limit the scope of US retaliation against the deviating country, requiring a higher discount factor to sustain cooperation. Once US tariffs exceed the Nash level, the US would not retaliate to any further deviation by other countries, and the required discount factor declines again.

A separate scenario examines what happens when other countries exclude the US from their tariff retaliation — as might happen, for instance, due to geopolitical considerations. In this case, China's minimum discount factor rises to 0.579, a more meaningful increase. The reason is that excluding the US from retaliation reduces both the gain from defection and the loss from a trade war, with the net effect of making cooperation harder to sustain.

Figure 1. Discount factor under partial exit of the US



From the figure, it displays the minimum discount factor under different initial US tariffs. (Higher discount factor = cooperation harder to sustain.)

How fragile is cooperation?

The minimum discount factors in our baseline are around 0.43 for China and the EU and 0.39 for the average country in the sample. Given a standard annual discount factor calibration in macro models of 0.96, cooperation appears straightforward to sustain. Several considerations, however, caution against this reading.

Policymakers may be effectively less patient than the textbook benchmark implies. Re-election risk can generate present-biased behaviour in trade policy (Amador 2003). Gabaix (2020) introduces the concept of cognitive discounting, under which limited foresight renders agents as-if less patient; the implied annual discount rate in that setting is near 0.5, only slightly above the minimum discount factors we recover. The duration of punishment also matters.


Figure 2 shows that with only one period of punishment before a return to cooperation, the required discount factor rises to approximately 0.75. Under grim-trigger strategies, cooperation is highly stable; under short punishment horizons, it is considerably more fragile.

Implications

The concern motivating much of the policy commentary since April 2025 has been that a US withdrawal from the cooperative trade regime could set off a cascade of protectionism among the remaining countries, as the terms of the cooperation game change for all parties.

Our results do not support this concern. The key determinants of cooperation sustainability — domestic and international trade shares — change too little following the US exit to materially alter the payoffs in the repeated game. This holds across a range of model specifications and parameter assumptions.

Figure 2. Discount factor and length of punishment



Notes: The figure displays the minimum discount factor required as a function of the number of punishment periods before reverting to the cooperation strategy.

Indeed, in the period since Liberation Day, global trade cooperation does not appear to be decreasing. Several significant trade agreements have been announced or concluded since April 2025: the EU-Mercosur agreement, the EU-India trade agreement, and the expansion of the Comprehensive and Progressive Agreement for Trans-Pacific Partnership (CPTPP).

If the US withdrawal activates this dynamic among remaining countries, it may generate more trade liberalisation rather than less. Our analysis of country incentives suggests the conditions for continued trade cooperation remain strong. ■

Barthélémy Bonadio is an Assistant Professor of Economics at NYU Abu Dhabi, Andrei Levchenko is the John W Sweetland Professor of International Economics at the University of Michigan, Ann Arbor, and Nitya Pandalai-Nayar is Assistant Professor of Economics at the University of Texas at Austin

Endnote

1. See the discussion in Ossa (2015).

References

- Amador, M (2003), "A Political Economy Model of Sovereign Debt Repayment", mimeo, Stanford University.
- Bagwell, K, CP Bown and RW Staiger (2016), "Is the WTO Passé?", *Journal of Economic Literature* 54(4): 1125–1231.
- Baldwin, R (1993), "A Domino Theory of Regionalism", NBER Working Paper 4465.
- Bonadio, B, Z Huo, AA Levchenko and N Pandalai-Nayar (2020), "[The role of global supply chains in the COVID-19 pandemic and beyond](#)", VoxEU.org, 25 May.
- Bonadio, B, Z Huo, AA Levchenko and N Pandalai-Nayar (2021), "Global Supply Chains in the Pandemic", *Journal of International Economics* 133: 103534.
- Bonadio, B, AA Levchenko and N Pandalai-Nayar (2026), "Falling Dominoes? The Impact of the US Exit from Free Trade on the Sustainability of Trade Cooperation", NBER Macroeconomics Annual.
- Costinot, A and A Rodríguez-Clare (2014), "Trade Theory with Numbers: Quantifying the Consequences of Globalization", *Handbook of International Economics* 4: 197–261.
- Gabaix, X (2020), "A Behavioral New Keynesian Model", *American Economic Review* 110(8): 2271–2327.
- Gros, D (1987), "A Note on the Optimal Tariff, Retaliation and the Welfare Loss from Tariff Wars in a Framework with Intra-Industry Trade", *Journal of International Economics* 23(3-4): 357–367.
- Horn, H, L Johannesson and PC Mavroidis (2011), "The WTO Dispute Settlement System 1995–2010: Some Descriptive Statistics", *Journal of World Trade* 45(6): 1107–1138.
- Huo, Z, AA Levchenko and N Pandalai-Nayar (2025), "International Comovement in the Global Production Network", *Review of Economic Studies* 92(1): 365–403.
- Ignatenko, A, A Lashkaripour, L Macedoni and I Simonovska (2025), "Making America great again? The economic impacts of Liberation Day tariffs", *Journal of International Economics* 157, 104138.

Johnson, HG (1951), "Optimum Welfare and Maximum Revenue Tariffs", *Review of Economic Studies* 19(1): 28–35.
Lashkaripour, A (2021), "The Cost of a Global Tariff War: A Sufficient Statistics Approach", *Journal of International Economics* 131, 103419.
Mei, Y (2020), "Sustainable Cooperation in International Trade: A Quantitative Analysis", *Journal of International Economics* 123, 103305.
Ossa, R (2014), "Trade Wars and Trade Talks with Data", *American Economic Review* 104(12): 4104–46.
Ossa, R (2015), "[WTO success: No trade agreement but no trade war](#)", VoxEU.org, 11 June.

This article was originally published on [VoxEU.org](#).

Europe's corporate tax trilemma

The EU's implementation of global anti-tax avoidance rules and the global minimum tax is widely regarded as a success. Madalena Barata da Rocha, Roel Dom, Pascal Saint-Amans and Bo Sangers consider the trilemma between tax sovereignty, revenue protection and investment neutrality that has emerged

Executive summary

The European Union's implementation of global anti-tax avoidance rules and the global minimum tax is widely regarded as a success. By substantially reducing opportunities for profit shifting, these reforms have strengthened revenue protection. But this success also has an underappreciated implication for the EU.

Profit shifting in part emerged as a response to differences in national corporate tax systems. It resulted in multinational firms reducing the effective tax burden associated with investing in higher-tax EU countries. It thus both eroded tax bases and offset the influence of tax differences on real investment decisions.

By constraining this form of 'self-help' while leaving the underlying diversity of national corporate tax systems largely intact, anti-avoidance reforms potentially increase the sensitivity of investment to tax differentials across the EU. Evidence shows that investment has become significantly more responsive to effective corporate tax rates in the last decade since coordinated anti-avoidance rules started to be discussed.

Reforms to tackle profit shifting have thus shifted the trade-off facing European policymakers to a trilemma involving tax sovereignty, revenue protection and investment neutrality. In this context – though not questioning the merits of combatting profit shifting – the success of anti-avoidance reforms strengthens the economic case for further coordination of corporate taxation within the single market, as anti-avoidance policy and tax harmonisation become increasingly complementary.

1 Introduction

A tension lies at the heart of the European Union's single market. Goods, services, capital and people move freely across borders, but corporate taxation remains organised largely along national lines. For companies, this creates complexity and ambiguity, increasing tax compliance costs, but also opportunities for tax avoidance, encouraging tax competition and distorting the allocation of investment.

Policymakers and academics have spent decades debating whether tax harmonisation could promote greater economic integration¹, yet progress towards harmonisation has been limited. Corporate taxation remains largely a national competence and EU-level tax initiatives have stalled repeatedly. Rather than reducing these differences directly through harmonisation, EU tax policy has thus focused increasingly on limiting the opportunities divergence creates for profit shifting from high to low-tax countries. In particular, over the last decade, reforms inspired by the Organisation for Economic Co-operation and Development/G20 Base Erosion and Profit Shifting (BEPS) project², culminating in the implementation of a minimum tax on the profits of large multinationals (known as Pillar Two), have substantially reduced opportunities for profit shifting.

Most analyses of profit shifting focus on revenue losses and fairness concerns. The BEPS reforms are considered a success because they help countries safeguard their tax revenues (eg. Contreras *et al* 2026). But these reforms also have an underappreciated consequence for economic efficiency. In a fragmented tax environment such as the EU, profit shifting reduces the extent to which investment decisions are driven by tax considerations rather than economic fundamentals.

The success of anti-avoidance reforms has therefore weakened a mechanism that allowed multinational firms to offset and reduce the effective tax burden associated with investing in higher-tax EU countries. Consequently, the economic importance of tax differentials between jurisdictions has increased again.

This changes the trade-offs facing European tax policy. A trilemma between tax sovereignty, revenue protection and investment neutrality has emerged. In the trilemma, if the goal for the EU is to combine revenue protection with economic efficiency, the price is a reduction in tax sovereignty: member states would need to accept greater convergence in tax rates and bases to prevent tax differentials from distorting investment. If, instead, tax sovereignty is to be preserved alongside revenue protection, economic efficiency will remain weakened.

Note, the argument is not that profit shifting should be tolerated. The adverse consequences it creates are real and the progress achieved through international tax cooperation should not be reversed. Rather, the consequence for the EU of progress on BEPS is that economic benefits of further tax harmonisation have increased. The success of anti-avoidance strengthens the case for continuing efforts to reduce differences in tax systems and to promote a more integrated corporate tax framework within the single market.

In this respect, the global minimum tax may offer more than a revenue floor: the common tax base it has already established for large multinational groups could itself serve as a building block for broader harmonisation, reducing the technical burden of constructing a unified EU corporate tax framework from scratch.

This *Policy Brief* proceeds in four steps. Section 2 briefly sets out the EU corporate tax framework, focusing on the contrast between limited harmonisation and the more successful implementation of common anti-avoidance rules. Section 3 develops the paper's central argument that restricting profit shifting can increase the sensitivity of real investment to corporate tax differentials and presents evidence from the EU single market alongside supporting evidence from other jurisdictions. Section 4 draws out the resulting policy trade-off between tax sovereignty, revenue protection and investment neutrality. Section 5 considers the implications for EU tax policy and sets out options for moving towards a more coherent balance between these objectives.

2 The EU corporate tax landscape

The EU has advanced substantially on coordinated anti-avoidance while leaving the underlying diversity of national corporate tax systems largely intact. This section briefly reviews both, since the interaction between them underpins the argument that follows.

2.1 Persistent fragmentation and limited positive harmonisation

Corporate income taxation across the EU is fragmented as it remains a national competence. EU legislation on direct taxation (ie. positive harmonisation)³ requires unanimity in the Council of the EU under Article 115 TFEU, meaning any EU country can veto a proposal⁴.

For the EU, if the goal is to combine revenue protection with economic efficiency, the price is a reduction in tax sovereignty

This is not without consequence as political preferences diverge: smaller, open economies tend to use competitive tax rates as instruments of their growth strategies, while larger economies with substantial domestic markets have less to gain from tax competition and more to lose from base erosion. Accordingly, they tend to favour harmonisation.

Tax systems are, moreover, deeply embedded in national social contracts, reflecting genuine differences in democratic preferences on how governments finance public services and redistribute income. As a result, ambitious EU proposals to harmonise to some degree corporate taxation have stalled repeatedly⁵. Progress has been confined to narrow directives that dismantle specific obstacles to crossborder activity.

These include the 2011 Parent-Subsidiary Directive (Council Directive 2011/96/EU), which eliminates double taxation of profits distributed between related companies in different member states, the 2009 Merger Directive (Council Directive 2009/133/EC – removes fiscal obstacles in crossborder reorganisations), the 2003 Interest and Royalties Directive (Council Directive 2003/49/EC – removed source taxation on intra-group interest and royalty payments), and the FASTER Directive of 2025 (Council Directive (EU) 2025/50 – seeks to simplify and accelerate withholding-tax relief procedures for crossborder investors).

2.2 Negative harmonisation and its limits

Judicial integration has gone further. The Court of Justice of the EU (CJEU) polices national tax rules in the context of the fundamental freedoms of the internal market, particularly the freedom of establishment and the free movement of capital. The CJEU has struck down measures that discriminate against crossborder activity.

But this ‘negative’ harmonisation cuts both ways. Because national tax measures that restrict free movement have traditionally been justified only in cases of “*wholly artificial arrangements*” (see Box 1), the ability of EU countries to defend their tax bases against intra-EU profit shifting has been constrained.

The 2006 Cadbury Schweppes ruling on the United Kingdom's controlled-foreign-company (CFC) rules provides a particularly clear illustration: EU countries may not apply CFC rules to subsidiaries that carry out real economic activity in other member states, even where those subsidiaries benefit from substantially lower tax rates (Box 1).

More recent case law suggests that the Court has become somewhat more permissive of abstract and general anti-abuse measures (Weber and Wattel 2025), but the underlying tension between market integration and base protection remains. Because the legislative route is so heavily obstructed, much of the work of integration has fallen instead to the CJEU.

2.3 Combatting base erosion and profit shifting (BEPS)

The corporate tax divergences that remain in place in the EU are not costless. They create opportunities for tax avoidance and incentivise tax competition. Reported profits are highly sensitive to tax-rate differences and are disproportionately booked in lower-tax member states (Huizinga and Laeven 2008; Beer *et al* 2018), with the resulting losses falling on higher-tax countries and flowing largely to other EU members (Nerudová *et al* 2023) (Figure 1).

Such competition has in turn pressed corporate tax rates downward, from above 40 percent in the mid-1980s to about 20–25 percent today, as EU countries adjust strategically to one another (Devereux *et al* 2008; Overesch and Rincke 2009).

In response, and in the wake of the 2008–2010 global financial crisis, when global revenue losses were estimated at \$100 billion to \$240 billion per year, the G20 and the OECD started the Base Erosion and Profit Shifting (BEPS) project in 2013 (OECD 2013; OECD 2015). EU countries, along with other BEPS members, have written common


Box 1. The Cadbury Schweppes ruling⁶

Cadbury Schweppes concerned the United Kingdom's controlled-foreign-company (CFC) rules, under which the UK tax authorities included the profits of subsidiaries established in low-tax jurisdictions (in this case, Ireland) in the tax base of resident parent companies. The objective was to prevent profit shifting and tax avoidance.

The CJEU held that such rules restrict the freedom of establishment and may only be justified where they target "*wholly artificial arrangements*", defined as subsidiaries lacking real substance and established solely for tax purposes. A genuine establishment (ie. one with real staff, premises and activity) therefore falls within the protection of the freedom of establishment; the mere fact that it benefits from a lower-tax regime does not make it a wholly artificial arrangement.

Consequently, EU countries may not apply CFC rules to subsidiaries established in other member states where those subsidiaries carry out real economic activities. This limits the ability of member states to tax profits shifted to lower-tax jurisdictions within the EU, while an equivalent domestic structure would have its profits taxed at the full national rate.

Figure 1. Corporate tax revenue lost to other jurisdictions due to profit shifting (% of GDP), 2022



Note: revenue winners are not shown; these are Ireland, Netherlands, Luxembourg, Belgium, Malta and Cyprus.
Source: Bruegel based on the [Atlas of Offshore World](#), EU Tax Observatory, and on World Bank GDP data.

these standards into national law to limit international tax avoidance. The EU has given effect to this agenda through three main instruments.

The 2016 Anti-Tax Avoidance Directive (ATAD, Council Directive (EU) 2016/1164), extended in 2017 by ATAD II (Council Directive (EU) 2017/952) established common minimum standards on interest deductibility, exit taxation, a general anti-abuse rule, hybrid mismatches and a CFC regime that attenuates the tension described in Box 1. As a minimum-standards instrument, it constrains avoidance without displacing national systems.

The Directive on Administrative Cooperation (DAC, Council Directive 2011/16/EU), originally adopted in 2011 and amended many times since, has built an increasingly comprehensive system of automatic information exchange between national tax authorities, attacking avoidance through transparency⁷.

The 2022 Minimum Tax Directive (Council Directive (EU) 2022/2523) occupies a somewhat special position. It transposed Pillar Two of the OECD/G20 global agreement into EU law, obliging member states to levy a top-up tax that ensures an effective rate of at least 15 percent on large multinational groups (those with consolidated revenues above €750 million)⁸.

By prescribing a minimum tax, Pillar Two blunts the advantage of shifting profits to low-tax jurisdictions and so targets tax competition itself. It leaves the broader diversity of national corporate tax systems intact. As such, it represents a significant, but partial, step towards corporate tax harmonisation.

The EU unanimity requirement on tax rules did not prevent agreement on ATAD, DAC or the minimum tax. Several factors account for the contrast. Anti-avoidance measures protect national tax bases rather than surrender them, so they align with the revenue interests of larger EU countries; they built on an externally negotiated OECD consensus,

which lowered the political cost of agreement and made resistance harder to justify; and even lower-tax states could accept a global minimum that bound their competitors as much as themselves.

The contrast between substantial EU progress on anti-avoidance and comparatively little progress on structural harmonisation matters because profit shifting and tax fragmentation are not independent phenomena. The former emerged, at least in part, as a response to the latter. Understanding how these two features interact is therefore essential to assess the broader consequences of anti-avoidance reforms.

3 The overlooked efficiency channel of profit shifting: loss of self-help

Standard analyses of profit shifting focus on its costs: it distorts profit allocation away from where economic activity takes place, erodes member state tax bases, generates costly planning structures that serve no productive purpose and creates competitive inequities between domestic firms (which cannot shift profits) and multinationals (which can). These distortions are real and well-documented. Globally, multinationals are estimated to shift a significant share of their profits to low-tax jurisdictions, at substantial cost to public revenues (Hines 2010; Tørsløv *et al* 2023; Guvenen *et al* 2022).

Yet this standard framing overlooks a less visible channel through which profit shifting interacts with real economic activity. By reducing the effective tax burden on investment in high-tax jurisdictions, profit shifting can offset to some degree the investment distortions that arise from tax differentials. In the EU, with 27 different statutory tax rates and a wide range of tax incentives, profit shifting has dampened the sensitivity of real investment to statutory tax differentials. Removing that ability to profit shift without addressing the underlying tax diversity may therefore increase the sensitivity of investment to tax differentials.

3.1 The investment channel: theory

The academic literature increasingly recognises that profit-shifting opportunities weaken the link between statutory corporate tax rates and real investment decisions (Dharmapala 2008; Hong and Smart 2010; Klemm and Liu 2019; Mongrain *et al* 2023). When multinationals can reallocate taxable profits across jurisdictions, the effective tax burden associated with investing in a high-tax country is reduced. Profit shifting thus lowers the effective cost of capital and can stimulate real investment even in relatively high-tax jurisdictions. The implication is that tax differentials exert a weaker influence on capital allocation when profit shifting is feasible, compared to when profit-shifting is not available.


Lax anti-avoidance regulation then functions as a substitute for statutory tax rate cuts by reducing the cost of capital for multinational firms. EU countries that tolerate profit shifting effectively operate corporate tax rates for multinationals that are below their headline rates.

Constraining profit shifting therefore changes the competitive equilibrium not only in terms of reported profits but also in terms of real investment incentives. Without compensating adjustments to rates or bases, anti-avoidance reform can therefore raise the effective tax burden on mobile investment, restoring the importance of tax differentials as a driver of location decisions.

3.2 Evidence from the EU single market and other jurisdictions

Intra-EU foreign direct investment data supports this supposition. Intra-EU FDI became measurably more tax-sensitive after the announcement of the BEPS initiative in 2015 (Figure 2). Before 2015, a one-percentage-point increase in the host-country effective average tax rate (EATR) was associated with an approximately eight percent decline in bilateral FDI flows. Since 2015, the same one-point increase has been associated with a 14 percent decline.

Figure 2. Sensitivity of FDI flows to effective tax rates, before and after BEPS



Notes: bars show the total implied effect of the EATR on inverse hyperbolic sine transformation of bilateral FDI flows (2013-2024). See the appendix for methodological details. Five non-OECD EU members (Bulgaria, Croatia, Cyprus, Malta and Romania) appear only as source countries.
Source: Bruegel.

Put differently, a given cross-country tax differential has exerted roughly 75 percent more influence over the location of real investment within the single market post-2015. This is consistent with profit-shifting constraints restoring the link between statutory tax burdens and genuine investment location decisions.

Cross-country evidence supports the evidence shown in Figure 2. Buettner *et al* (2018), using a worldwide panel of multinational affiliates, found that introducing or tightening a thin-capitalisation rule increased the tax-rate sensitivity of FDI in high-tax host countries. De Mooij and Liu (2021) corroborated this on the real investment margin.

Exploiting variation in the tightness of thin-capitalisation safe-haven ratios across 34 countries between 2006 and 2014, they found that tighter rules significantly reduced real capital accumulation by affected multinationals, with the adverse effect increasing in line with the strictness of the rule and the level of the host-country tax rate.

Evidence from case studies further supports this pattern. In a study of the impact of the repeal of Section 936 of the US Internal Revenue Code – a provision that allowed US multinationals to shift profits to Puerto Rican affiliates – Suárez Serrato (2018) found that eliminating this shifting channel raised the effective tax cost of domestic investment, with measurable real consequences.

Affected firms substantially reduced domestic investment and domestic employment. The effects were amplified at the local level, with labour markets exposed to profit-shifting firms experiencing declines in wages, employment and property values, while reliance on government transfers increased.

A closer analogue to the EU experience comes from Bilicka *et al* (2022) who exploited the United Kingdom's 2010 Worldwide Debt Cap (an anti-avoidance rule limiting interest deductibility relative to a firm's global debt) as a quasi-

natural experiment. Affected multinationals reduced their UK debt and increased debt held in foreign affiliates, but critically also reallocated some of their real operations away from the UK. This showed that anti-avoidance reforms do not simply curtail paper profit shifting but can trigger real activity relocation, with direct consequences for domestic investment and employment.

Bilicka *et al* (2024) provided some of the most direct structural evidence on the investment consequences of restricting profit shifting. Based on UK administrative tax return data, they found that a one percentage point rise in the host-country corporate tax rate reduces capital accumulation by 0.8 percent when profit shifting is available, but by 1.8 percent in its absence. In other words, profit shifting dampens the investment sensitivity of multinationals to tax rate changes by more than half.

To be clear: none of this implies that profit shifting should be tolerated. The argument is narrower. In the fiscally fragmented EU environment, profit shifting made multinationals less sensitive to tax differentials, functioning as an imperfect substitute for the tax harmonisation that political constraints have so far prevented.

Anti-avoidance reforms have changed that. Access to consumers, labour quality, infrastructure, legal certainty and other market fundamentals continue to anchor investment decisions, and the EU single market remains competitive in these respects. But at the margin, tightening of anti-avoidance rules has increased the importance of remaining tax differentials across EU countries.

4 The emerging trilemma

Debates on European corporate taxation have long been framed as a tension between tax sovereignty and economic efficiency (the latter understood here simply as investment neutrality)⁹. Greater tax harmonisation would

reduce tax-induced distortions within the single market but would constrain the ability of member states to design their own tax systems.

Preserving national sovereignty, conversely, implies differences in tax rates and bases that distort the location of investment. This framing, however, is incomplete. Recent developments suggest that a third objective – revenue protection – must be brought into the picture. The interaction between all three generates a trilemma: no configuration secures all three at once.


Tax sovereignty sits at the apex of the trilemma (Figure 3), with revenue protection and economic efficiency at the two base corners. The trade-offs between them are real. Preserving tax sovereignty (ie. the ability of EU countries to set their own rates, define their own bases and deploy taxation as an instrument of economic policy) produces the kind of differentiation that distorts multinationals' investment decisions and creates incentives to shift profits across borders.

Pursuing revenue protection constrains those profit-shifting opportunities but also dampens the mechanism that enables multinationals to reduce the effective tax burden associated with investing in higher-tax EU countries, increasing the tax sensitivity of investment.

Finally, moving to location-neutral investment allocation requires member states to surrender some of the fiscal autonomy through which national preferences over the tax-and-spend mix are expressed.

The pre-BEPS equilibrium occupied a position towards the right of the triangle, closer to investment neutrality and tax sovereignty, but at a cost to revenue protection. The negative harmonisation achieved through the CJEU reinforced this dynamic: rulings such as Cadbury Schweppes (Box 1) constrained EU countries from applying CFC

Figure 3. An emerging trilemma



Source: Bruegel.

rules to subsidiaries with genuine economic substance, limiting their ability to tax profits shifted to lower-tax jurisdictions within the EU.

Profits were systematically booked away from where economic activity occurred, with the tax revenue losses falling on higher-tax countries. But this system meant firms could make real investments in those high-tax countries that may not have been worthwhile in the absence of profit shifting.

The BEPS reforms from 2015 on have shifted the equilibrium towards the revenue protection side. BEPS-linked EU policy measures such as ATAD and DAC have achieved what structural tax harmonisation attempts have failed to deliver: binding, enforceable common rules that have substantially curtailed profit-shifting opportunities across the single market. But this comes at a cost to economic efficiency, making capital flows more sensitive to tax differentials again (section 3), meaning that investment within the single market is increasingly shaped by differences in national tax systems rather than by underlying economic fundamentals.

The post-2015 position therefore sits closer to the revenue protection and tax sovereignty vertices but has moved the single market further from the investment neutrality corner than before.

The question of the EU's objective follows directly. If the goal is to combine revenue protection with economic efficiency, the price is a reduction in tax sovereignty: member states would need to accept greater convergence in tax rates and bases in order to prevent tax differentials from distorting investment.

If, instead, tax sovereignty is to be preserved alongside revenue protection, investment neutrality will remain compromised, and the distortions documented in section 3 will persist or intensify. The trilemma cannot be dissolved; it can only be managed by making the trade-off explicit and choosing which objectives to prioritise.

The EU implementation of Pillar 2 has introduced an element of coordination by establishing a common minimum effective tax rate of 15 percent for in-scope multinational enterprises. The 15 percent global minimum tax (GMT) may be seen as a step towards combining revenue collection, by putting a floor under tax competition, and investment neutrality. The GMT partly reduces competition on rates, to the extent that effective rates below 15 percent were available in the EU (and elsewhere).

Early evidence suggests that the GMT has not had an impact on investment (Contreras *et al* 2026). Although this does not amount to full rate harmonisation, since member states remain free to set their own statutory corporate tax rates, it does limit tax competition below the effective minimum.

In addition, several elements of a package issued by the European Commission in June 2026 are relevant in light of these trade-offs¹⁰. These include proposals to extend withholding tax exemptions on intra-EU payments of dividends, interest and royalties, aligning the interest limitation rules under the Anti-Tax Avoidance Directive, broadening the exchange of information on ownership of immovable property and introducing a common minimum standard for the immediate expensing of R&D expenditure.

A proposed safeguard against double non-taxation shares some features with the idea of external tax borders. While falling well short of full corporate tax harmonisation, the package suggests a shift to combining anti-avoidance measures with greater coordination of tax rules.

To the extent that this reduces both opportunities for profit shifting and tax-related frictions within the single market, it may help reconcile the objectives of revenue protection and economic efficiency. Whether the package ultimately achieves this balance, however, remains an empirical question that can only be assessed once the reforms have been implemented.

The current trajectory of EU tax policy has therefore not fully resolved the trilemma. Anti-avoidance reforms have protected revenue while, despite some coordination, tax sovereignty is mostly preserved, at the cost of economic efficiency. Section 5 considers what a more deliberate response might look like.

5 Policy recommendations

Recent EU tax policy has been largely shaped by a focus on protecting the revenue base by limiting profit shifting, with little consideration of measures to limit some of the potentially negative spillover effects to the real economy. This view is incomplete and does not account for complex trade-offs. Anti-avoidance reforms have curtailed profit shifting and protected member states' tax bases, but they have also removed a mechanism that, unintentionally, has cushioned real investment from the EU's underlying tax diversity.

Ignoring this trade-off would lead policymakers to treat further anti-avoidance tightening as unambiguously beneficial, when in fact it involves real and rising efficiency costs as remaining tax differentials regain influence over investment decisions. None of this should be read as an argument against anti-avoidance enforcement. But good policy design requires the trade-offs to be acknowledged and dealt with, rather than ignored.

- **If investment-neutrality is preferred, then corporate income tax harmonisation should be stepped up.** Reducing distortions to investment within the EU requires a move away from national tax sovereignty, which has proved hard so far, with very limited direct harmonisation. EU countries view common tax rules not only as a direct threat to their prerogatives but also a route to greater CJEU intervention in their tax affairs.
- **The current geopolitical environment justifies a more ambitious approach to harmonisation.** Corporate income tax accounts for a relatively modest share of total government revenue (about 10 percent of revenue in the EU on average¹¹), suggesting that the revenue stakes of harmonisation, while real, should not be

overstated, and that scope remains for cross-country differences in tax-to-GDP ratios, reflecting national social preferences. Smaller countries, which currently collect shares of EU-wide corporate income tax revenue that are disproportionate to their economic sizes, would likely see this revenue fall in a harmonised system.

However, part of this revenue advantage already reflects profit-shifting flows that anti-avoidance reforms are progressively eliminating, and the additional top-up tax from the global minimum tax is temporary and will ultimately shrink. A coordinated approach, with the possibility of temporary compensation, would be better at the EU level.

- **The OECD/G20 Pillar Two framework provides a foundation for a harmonised tax base.** Some of the technical foundations for deeper harmonisation are already in place, making progress more attainable than one might think. The EU has transposed the Pillar Two framework into binding law through the Minimum Tax Directive (Council Directive (EU) 2022/2523). Under these rules, multinational enterprises must calculate effective tax rates using a common set of accounting and tax-base rules derived from consolidated financial accounts. A significant part of the complexity associated with establishing a common tax base has therefore been addressed and could be drawn on directly as a foundation for broader harmonisation.

While the minimum-tax rules are widely recognised as extremely complex, large companies have already absorbed this complexity. They could apply the same common tax-base calculations to a broader harmonised regime at limited additional cost. This favours an approach that targets large firms first: the segment of the economy most exposed to profit shifting. Smaller firms could continue to operate under national systems.

- **The EU should agree a common effective tax rate, or at least a more coordinated rate structure.** Rather than treating this framework as an additional layer on top of existing national systems, the EU could use it as

the foundation for a more comprehensive harmonised corporate tax regime for large companies, specifically for large multinationals within the scope of Pillar 2. This would target about half of corporate income tax revenue as these companies contributed an average of 47.1 percent of total corporate tax revenues in 2022 (OECD 2025).

- **If such a comprehensive reform proves unattainable in the near term, the EU should seek to reduce the frictions created by the current fragmented system.** As a first step, the European Commission could systematically map the remaining tax frictions that continue to impede crossborder activity within the single market.

Remaining withholding taxes on crossborder payments within the single market should be eliminated. Existing directives already provide exemptions in specific cases¹², but broader removal of withholding taxes could reduce barriers to crossborder investment and improve capital allocation across member states.

- **Internal harmonisation needs to be accompanied by an external corporate income tax border.** Greater internal integration within the single market could increase the incentive for firms to move the tax base outside the EU altogether, rather than between member states. This risk should be addressed alongside any move towards internal harmonisation, not subsequent to it.

This could be achieved through an EU-wide approach to an external border of corporate income taxation, through consistent and coherent withholding taxes, as currently provided for in bilateral tax treaties in an uncoordinated manner. This could help safeguard member states' tax bases, while preserving the benefits of a more integrated internal market.

- **Finally, the EU should enhance competitiveness by coordinating tax-incentive policies that pursue common policy objectives.** EU countries increasingly use targeted tax provisions to promote activities such as green investment, innovation, strategic industries and technological development. Though these objectives are often shared across the EU, uncoordinated national incentives can generate subsidy races and distort investment decisions.

The European Commission (2025b) has issued recommendations establishing common principles for the design of clean-industry incentives, such as accelerated depreciation and tax credits for clean investment, but they are non-binding, leave the generosity of incentives to national discretion and are confined to clean-tech and decarbonisation investment.

In other words, they standardise the form of incentives without constraining the intensity of competition between member states offering them. A more effective approach would require binding limits on incentive generosity, not merely common design principles, and would need to extend coverage beyond the clean-tech domain to other contested categories of mobile investment. ■

ABOUT THE AUTHORS

Madalena Barata da Rocha is a Research Analyst, Roel Dom is a Research Fellow and Senior Researcher at EU Tax Compass, Pascal Saint-Amans is a Senior Fellow and Project Lead at EU Tax Compass, and Bo Sangers is a Research Assistant, all at Bruegel

Endnotes

1. See, for example Kopits (1992) for an early contribution to the long-standing debate on tax harmonisation in the context of European economic integration, and Andersson et al (2025) for a more recent study on the economic and administrative costs of fragmented tax rules in the EU.
2. See OECD, '[Base erosion and profit shifting \(BEPS\)](#)'.
3. EU tax integration is conventionally understood through the distinction between negative and positive harmonisation. Negative harmonisation is the removal of national rules that obstruct the internal market, achieved through judicial review rather than legislation. Positive harmonisation is the construction of common rules through EU legislation.
4. Article 115 also limits the depth of harmonisation achievable: it empowers the Council only to "issue directives", whereas Article 113 authorises 'provisions', a broader term that can also support regulations, or EU laws applied directly and uniformly. Since directives leave implementation to member states, direct-tax measures can only approximate national systems, never unify them through directly applicable rules. Direct-tax harmonisation is thus doubly constrained: by unanimity and by the instrument itself.
5. Major initiatives include the Common Consolidated Corporate Tax Base (CCCTB), proposed in 2011 and relaunched in 2016, which sought a single set of rules for computing and apportioning corporate profits, but failed. The idea was revived in 2023 under the Commission's BEFIT (Business in Europe: Framework for Income Taxation) proposal (European Commission 2023). However, BEFIT remains pending. See the European Commission's 2026 work programme (European Commission 2025a).
6. See CJEU ruling, '[Cadbury Schweppes plc and Cadbury Schweppes Overseas Ltd vs. Commission](#)', C-196/04, of 12 September 2006.
7. The successive DACs cover financial account data (DAC2), advance crossborder tax rulings (DAC3), country-by-country reports (DAC4), reportable crossborder arrangements (DAC6) and, more recently, income earned through digital platforms (DAC7) and crypto-assets (DAC8).

8. Pillar Two is part of the OECD/G20 two-pillar deal agreed in October 2021. Whereas Pillar Two sets a floor on the effective tax rate paid by large multinationals, Pillar One creates a new taxing right for market jurisdictions. Unlike Pillar Two, Pillar One has not been implemented because of a lack of global consensus.
9. The efficiency referred to here is narrow: neutrality in the location of real investment, not efficiency in the broader welfare sense. Profit shifting is not costless. It distorts organisational form, since the planning opportunities it relies on are available to multinational groups but not to purely domestic firms; it distorts financing and the location of intangibles; and it absorbs real resources in tax planning and enforcement. The pre-interventions framework was not neutral for investment let alone efficient in any broader sense.
10. See European Commission news article of 24 June 2026, [‘European Commission proposes landmark tax simplification package to streamline compliance and boost competitiveness’](#).
11. Eurostat, [‘Main national accounts tax aggregates’](#). Missing data for Germany, Spain and Hungary.
12. Such as the Parent-Subsidiary Directive (Council Directive 2011/96/EU) and the Interest and Royalties Directive (Council Directive 2003/49/EC).
13. See ZEW, [‘Mannheim Tax Index’](#).

References

- Andersson F, I Begg, T Bührle, A Korenbilt, K Patazatou and A Thomadakis (2025) [The future of EU tax policy harmonisation – Cost of non-Europe Report](#), European Parliamentary Research Service.
- Beer, S, R de Mooij and L Liu (2019) [‘International Corporate Tax Avoidance: A Review of the Channels, Magnitudes, and Blind Spots’](#), Working Paper 2018/168, International Monetary Fund.
- Bilicka, KA, MP Devereux and İ Güçeri (2024) [‘Tax policy, investment and profit shifting’](#), NBER Working Paper 33132, National Bureau of Economic Research.

- Bilicka, KA, Y Qi and J Xing (2022) *'Real responses to anti-tax avoidance: Evidence from the UK Worldwide Debt Cap'*, *Journal of Public Economics* 214: 104742.
- Buettner, T, M Overesch and G Wamser (2018) *'Anti profit-shifting rules and foreign direct investment'*, *International Tax and Public Finance* 25(3): 553-580.
- Contreras, L, F Hugger and T Zawisza (2026) *'MNE Responses to the Global Minimum Tax'*, OECD Taxation Working Papers 77, Organisation for Economic Co-operation and Development.
- De Mooij, R, & Liu, L (2021) *'At a cost: The real effects of thin capitalization rules'*, *Economics Letters* 200: 109745.
- Devereux, MP, B Lockwood and M Redoano (2008) *'Do countries compete over corporate tax rates?'* *Journal of Public Economics* 92(5-6): 1210-1235.
- Dharmapala, D (2008) *'What problems and opportunities are created by tax havens?'* *Oxford Review of Economic Policy* 24(4): 661-679.
- European Commission (2025a) *'Commission work programme 2026'*, COM(2025) 870 final.
- European Commission (2025b) *'Tax incentives to support the Clean Industrial Deal and in light of the Clean Industrial Deal State aid Framework'*, Recommendation (EU) 2025/1307.
- European Commission (2023) *'Proposal for a Council Directive on Business in Europe: Framework for Income Taxation (BEFIT)'*, COM(2023) 532 final.
- Guvenen, F, RJ Mataloni Jr, DG Rassier, and KJ Ruhl (2022) *'Offshore profit shifting and aggregate measurement: Balance of payments, foreign investment, productivity, and the labor share'*, *American Economic Review* 112(6): 1848-1884.
- Hong, Q and M Smart (2010) *'In praise of tax havens: International tax planning and foreign direct investment'*, *European Economic Review* 54(1): 82-95.
- Huizinga, H and L Laeven (2008) *'International profit shifting within multinationals: A multi-country perspective'*, *Journal of Public Economics* 92(5-6): 1164-1182.
- Hines, JR (2010) *'Treasure islands'*, *Journal of Economic Perspectives* 24(4): 103-126.

- Klemm, A and L Liu (2019) *'The Impact of Profit Shifting on Economic Activity and Tax Competition'*, Working Paper 2019/287, International Monetary Fund.
- Kopits, G (1992) *Tax Harmonization in the European Community: Policy Issues and Analysis*, International Monetary Fund.
- Mongrain, S, D Oh and T Van Ypersele (2023) *'Tax competition in the presence of profit shifting'*, Journal of Public Economics 224: 104940.
- Nerudová, D, M Dobranschi, V Solilová and M Litzman (2023) *'Onshore and offshore profit shifting and tax revenue losses in the European Union'*, Economic Modelling 119: 106111.
- OECD (2013) *Action Plan on Base Erosion and Profit Shifting*, Organisation for Economic Co-operation and Development.
- OECD (2015) *BEPS Project Explanatory Statement: 2015 Final Reports*, Organisation for Economic Co-operation and Development.
- OECD (2025) *Corporate Tax Statistics 2025*, Organisation for Economic Co-operation and Development.
- Overesch, M and J Rincke (2009) *'What Drives Corporate Tax Rates Down? A Reassessment of Globalization, Tax Competition, and Dynamic Adjustment to Shocks'*, CESifo Working Paper No. 2535, CESifo.
- Suárez Serrato, JC (2018) *'Unintended consequences of eliminating tax havens'*, NBER Working Paper 24850, National Bureau of Economic Research.
- Tørsløv, T, L Wier and G Zucman (2023) *'The missing profits of nations'*, The Review of Economic Studies 90(3): 1499-1534.
- Weber, D and P Wattel (2025) *'Free movement and tax base integrity'*, in S Douma, O Marres, H Vermeulen and D Weber (eds) *European Tax Law Volume I: General Topics and Direct Taxation* 9th ed, Wolters Kluwer.

This Policy Brief is based on research funded by the Gates Foundation. The findings and conclusions contained within are those of the authors and do not necessarily reflect positions or policies of the Gates Foundation. We thank Vitor Gaspar, Ruth Mason, Hans Geeroms for their comments, and participants in the Bruegel Research Meeting. This article is based on [Bruegel Policy Brief](#) Issue n°17/26 | August 2026.

Appendix

This appendix summarises the methodology, data and results of the regression analysis assessing the tax-sensitivity estimates of foreign direct investment (FDI), presented in section 3.2 and in Figure 2.

Research question. The analysis asks whether bilateral FDI flows became more sensitive to host-country corporate taxation after the 2015 OECD/G20 BEPS package curtailed profit-shifting opportunities. The underlying logic is that restrictions on profit shifting remove a channel that previously allowed multinationals to separate the location of reported profit from the location of real investment; if so, the location of real investment itself should become more responsive to tax differentials once that channel narrows.

Empirical specification. The model is a panel gravity regression of bilateral FDI flows:

$$\text{IHS}(\text{FDI}_{ijt}) = \beta_1 \text{TaxRate}_{jt} + \beta_2 \text{TaxRate}_{jt} \times \text{Post_BEPS}_t + \beta X_{(i)jt} + \alpha_{ij} + \lambda_t + \mu_i \times \lambda_t + \varepsilon_{ijt}$$

Where FDI_{ijt} is the flow from source country i to host country j in year t , transformed using the inverse hyperbolic sine transformation (which behaves like a log transformation while retaining zero and negative flows). TaxRate_{jt} is the host-country effective average tax rate; Post_BEPS_t equals one from 2015 onward. The coefficient of interest is β_2 . The pre-BEPS tax sensitivity of FDI is given by β_1 , and the post-BEPS sensitivity by $\beta_1 + \beta_2$. So, a negative and significant β_2 indicates that FDI became more responsive to host-country taxation after BEPS.

The specification includes three sets of fixed effects. Country-pair fixed effects (α_{ij}) absorb stable bilateral characteristics (eg. distance, language, trade and treaty history). Year fixed effects (λ_t) absorb global shocks common to all pairs in a given year (worldwide business cycles, global financial conditions, broad trends in corporate income tax rates). Source-country-by-year fixed effects ($\mu_i \times \lambda_t$) absorb time-varying conditions specific to the investing

country (eg. domestic tax reforms, financial constraints, surges or slumps in outward investment) so that the remaining variation isolates how the host country's tax rate affects the destination of a given source country's capital, net of what is happening in the source country itself. $X_{(i)jt}$ is a vector of time-varying control variables, including bilateral trade and time-varying host-country characteristics commonly used in the FDI gravity literature. The host-country controls account for observable determinants of FDI that vary over time within destination countries, helping to isolate the effect of corporate taxation on investment location decisions.

Data. The sample covers bilateral FDI flows among all EU27 country pairs, 2013-2024, with the host country as the reference unit and the partner country as the counterpart. FDI flows are drawn from the OECD bilateral FDI database, measured on a directional basis in current US dollars and converted to euros using annual OECD exchange rates; Special Purpose Entities (SPEs) are excluded, since SPE-related flows are driven primarily by tax and intra-group financing considerations rather than productive investment. Coverage gaps mean that Bulgaria, Croatia, Cyprus, Malta and Romania appear only as source (not host) economies; the resulting panel thus covers 22 host and 27 source countries.

The main explanatory variable, the effective average tax rate (EATR), is drawn from the ZEW Effective Tax Levels database (Mannheim Tax Index Update 2025¹³) and captures the tax burden on a profitable investment, incorporating the statutory rate and base-defining features such as depreciation allowances. Host-country controls include GDP (Eurostat, log), bilateral imports from the source into the host country (Eurostat, log, capturing trade integration), labour productivity (GDP per hour worked at PPP, OECD, log), institutional quality (the World Bank's rule-of-law indicator, range -2.5 to 2.5) and general government expenditure (Eurostat, log).

Table A1. Regression results (post-BEPS break: 2015)

	(1)	(2)	(3)	(4)
$EATR_{Host}$	-0.076*	-0.0696*	-0.0918*	-0.0825*
	(0.0405)	(0.0394)	(0.0478)	(0.0485)
$EATR_{Host} \times \text{Post BEPS}$	-0.0461*	-0.0485**	-0.0599**	-0.0599**
	(0.0246)	(0.0245)	(0.0289)	(0.0289)
Controls				
$GDP_{Host} \text{ (log)}$			-3.622*	-4.853**
			(1.856)	(2.276)
$Imports_{Source, Host} \text{ (log)}$			0.542*	0.566**
			(0.279)	(0.278)
$Labour \text{ productivity}_{Host} \text{ (log)}$			4.664	5.361*
			(2.876)	(2.899)
$Rule \text{ of law}_{Host}$				-1.124*
				(0.671)
$Public \text{ spending}_{Host} \text{ (log)}$				1.341
				(1.428)

Model information				
Pair FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Source x Year FE	No	Yes	Yes	Yes
Observations	6208	6208	5862	5862
R-squared	0.168	0.216	0.219	0.219

*Notes: Dependent variable: inverse hyperbolic sine transformation of net bilateral FDI flows in million euros, non-SPE, BMD4. Estimator: Panel OLS with two-way fixed effects (country pair and year in all specifications) and source country x year fixed effects in specifications 2, 3 and 4. Standard errors clustered at country-pair level in parentheses. EATR: corporate-level effective average tax rate, ZEW Mannheim Tax Index. post-BEPS = 1 if year > 2015. Five non-OECD EU members (BGR, HRV, CYP, MLT, ROU) appear only as source countries. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.*

Source: Bruegel.

Europe's federated democratic learning state

Europe is moving towards invoice-level electronic VAT reporting. Thiemo Fetzner argues that, with AI-assisted processing of the information contained in electronic VAT invoices, the tax system can become not just a compliance and revenue-raising instrument, but a protected intelligence map of the economy

Innovations like electronic payments and e-invoicing are making it much easier to trace economic activity (OECD 2024). At the same time, AI-assisted knowledge retrieval is increasingly substituting humans in a broad array of information intermediation tasks. Taken together, these developments can shift the informational boundaries of the state: what the state can observe, the knowledge it can derive, and how it may deploy that knowledge in the management of the economy.

The US and China have responded to the digitisation of economies and societies in radically different ways. In the American model, private players intermediate the economy's data trace, causing an extreme concentration of power. Payment service providers, for example, benefit from large network externalities, steady intermediation fees from cash flows, and information rents that accrue outside the ordinary social contract due to widespread profit shifting. The same applies to search, compute, advertising, platform work, and digital identity.

In the Chinese model, the state retains much more of that capacity directly, but at the cost of political hierarchy and coercion. Europe needs digital state capacity without that societal form, and digital sovereignty without platform dependence. By 2050, Europe should have built an alternative model: a federated democratic learning state that leans on public-data-derived evolving knowledge graphs and AI as an accountability and coordination layer with built-in privacy.

Digitisation of tax systems to build a new informational architecture

The unsung hero through which this may become concrete is the electronic VAT invoice. Europe and much of the G20 – except the US – is moving towards invoice-level VAT. This means every firm-to-firm and firm-to-consumer economic transaction becomes technically legible to the tax system in near real time. If every invoice contains structured or semi-structured line items, AI-assisted processing can make the tax system more than a compliance and revenue-raising instrument.

Such a system can become a protected intelligence map of the economy: who trades what with whom, which capabilities exist, where, and who produces what value added. A 'VAT firm-to-firm transaction network' represents a knowledge graph that measures the granular economic ties between essentially all economic units and agents, providing knowledge of what firms can produce through linking the observed outputs and measured inputs with a production-network view.

For Europe, a unique opportunity emerges to evolve the Single Market to become an intelligence and orchestration platform of a modern state that leans on real-time, high-authority administrative data

Granular network measurement is key to understanding aggregate resilience and the propagation of economic shocks (Acemoglu *et al* 2012). AI is already helping shape even more granular production-network representations of the tangible economy (Fetzer *et al* 2024b).

Putting this public-stewarded VAT-derived knowledge graph, combined with AI, to work can fundamentally reshape the role of the state and hence, redefine the boundaries of the firm. For Europe, a unique opportunity emerges to evolve the Single Market to become an intelligence and orchestration platform of a modern state that leans on real-time, high-authority administrative data.

AI is integral to ensure that underlying data can be operationalised in a way that is compatible with strict privacy, to protect sensitive data, and to maintain a market-based economic organisation. It can ensure that Europe moves to being governed by (trusted) statistics rather than stories. Present discussions on the future of the EU, formulated by landmark reports by Draghi (2024) and Letta (2024), have not emphasised the importance of the information architecture that may make these agendas administratively feasible.

I sketch one application in some detail: public procurement. Similar arguments can be made for innovation policy, technology transfer, and many other areas of public good provision.

Procurement as an example to incorporate knowledge graphs

Public procurement is often slow and inefficient. Bureaucracy, process, and due diligence are costly by-products of past corruption episodes. Even where there is no favouritism, its mere possibility can generate stories, litigation risk, and cycles of distrust.

The administrative state's response to any (perceived) scandal is more bureaucracy, raising participation costs, with the side effect of skewing public contracts to larger firms that have the long breadth alongside the required legal and administrative capacity.

A knowledge graph of the physical and material world, combined with an interoperable and digitally legible administrative state, consisting of company registries, and a capability graph anchored on production network knowledge fed with granular invoice-level VAT data, can vastly increase public procurement effectiveness. They can make the flow of physical goods through the economy visible in a way that Bill Phillips could only have dreamed of when designing his hydraulic computer of macroeconomic flows.

The process could look as follows: if the state wants to buy a good or service, it can use protected VAT-derived knowledge to suggest and invite a set of potential suppliers that – based on their economic classification, the value added, the observed inputs, and outputs – may possess the capability to provide that good or service.

AI-assisted retrieval is essential as it ensures that no human official can browse the underlying sensitive raw data that is built on production-network knowledge and VAT-transaction data. AI doesn't magically guarantee neutrality; rather, it makes the process auditable, query-bound, and outsize influence-blind in a way that can be regularly and automatically audited.

This changes the role of public authority. Industrial policy, for example, as often implemented through procurement, becomes less a matter of discretionary selection and more a process of structured self-learning and discovery based on broad, rule-based access across a range of tournaments. This reduces the risk of rentierism through 'revolving doors' and investments in proximity to power.

Once capability-eligible suppliers have been identified, access to bidding opportunities can be opened for a (stratified) random sample. Firms that are not registered and licensed, use coercive labour market practices, or infract in other ways on the social contract could be excluded from participation, creating an indirect demand for formalisation. It may also indirectly help to level the playing field.

By incorporating randomisation, the political economy of state demand changes. Firms can spend less time lobbying for contracts and focus more on delivery. Business development costs decrease. Further, if firms enter an auditable stochastic process, then not being selected becomes closer to an idiosyncratic risk, against which firms may, in equilibrium, find it cheaper to borrow, compared to a world where access to finance is also politicised.

In such a world, the policymaker's role shifts from managing to overseeing a near automated process. Procurement outcomes are shaped through sample stratification by factors such as region, firm size, labour standards, employment-structure, or firm-age. Each of these dimensions may reflect political decisions, but corruptive self-selection has been neutralised.

In essence, the policymaker's role is to develop the prompt for an AI agent to identify candidate suppliers from the publicly stewarded knowledge graph. Stochastic fairness can be guaranteed and audited, and – thanks to the steady flow of data – the underlying data evolves with the economy and, as a result, so does the knowledge graph.

Privacy, democratic accountability, and trust

A European learning state cannot be built on raw-data exposure but must guarantee privacy. Citizens and firms will only support a pivot to extensive, public-stewarded data if they know that data shared with public bodies will only be used for lawful, bounded, transparent purposes with efficiently challenging processes. This is the core concern in work on the informational boundaries of the state (Fetzer *et al* 2024a).

Public institutions should be able to fully utilise the informational resources they acquire, computing more while seeing less. A robust public data architecture paired with AI for trusted private knowledge retrieval is precisely the combination that can move societies towards an equilibrium with maximal knowledge sharing, minimal data exposure, and, crucially, without running the risk of turning into a Chinese- or American-style equilibrium.

Having transparency and auditability over how underlying data are being used is one pillar that can ensure that economic agents supply accurate high-quality data. The VAT has a built-in monitoring logic because transactions are relational. A reported sale is someone else's reported purchase. Dual reporting makes manipulation more costly, as false reporting requires coordination.

The threat of (automatic) exclusion, for instance from access to public demand, can become an additional disciplining device. None of this removes the need for audit, but it means the system is not based simply on self-reported claims; it is a network of mutually constraining records. Naturally, audit costs can be significantly lower in such a setting.

The development of a modern learning state can improve public accountability and repair democratic trust in a world of perceived polarisation. As rewards to proximity to government decline, the returns to lobbying fall. This can recalibrate the firm-size distribution by improving small and medium-sized companies' access to public sector demand.

Faster procurement can improve performative state capacity. Auditable, human-not-in-the-loop retrieval processes can deprive malign actors of narrative ammunition around non-transparent allocation. Fewer resources are wasted on pork-barrel politics and defensive compliance.

Most importantly, policymaking becomes more future-oriented as the short-term returns to using procurement to divert public resources decline. Political competition can then move towards the battle over ideas: what visions for individual and collective wellbeing should Europe build? ■

Thiemo Fetzner is Professor of Economics at Universität Bonn; Theme Leader for Globalisation and Global Crises at the Centre for Competitive Advantage in the Global Economy (CAGE); and Professor of Economics at the University of Warwick

References

- Acemoglu, D, VM Carvalho, A Ozdaglar, and A Tahbaz-Salehi (2012), "The network origins of aggregate fluctuations", *Econometrica* 80(5): 1977–2016.
- Draghi, M (2024), *The future of European competitiveness*, European Commission.
- Fetzer, T, C Shaw, and J Edenhofer (2024a), "[Informational boundaries of the state](#)", CEPR Discussion Paper 18773.
- Fetzer, T, PJ Lambert, B Feld, and P Garg (2024b), "[AI-generated production networks: Measurement and applications to global trade](#)", CEPR Discussion Paper 19708.
- Letta, E (2024), *Much more than a market*.
- OECD (2024), "Digital public infrastructure for digital governments", OECD Public Governance Policy Papers No. 68.

Author's note: I would like to thank Olivier Blanchard, Jakob Schneebacher and Jacob Edenhofer who provided helpful comments that significantly helped streamline the exposition. This article was originally published on [VoxEU.org](#).