

Europe's AI Regulation Problem: Innovation, Growth, and the Cost of Compliance

Advanced Macroeconomics

Amsterdam University College

Adrian Alvarez Sande

May 2026

Summary

The EU AI Act is the first comprehensive legal framework for artificial intelligence in the world. Its goals, including safety, accountability, and transparency, are reasonable ones. This brief does not argue against regulating AI. What it argues is that the current design imposes compliance costs that are not well-calibrated to firm size, risks slowing long-run productivity growth, and produces macroeconomic consequences that the policy debate has largely ignored. Drawing on the Solow growth framework, the government intertemporal budget constraint, and the Taylor Rule, the brief shows that excessive AI regulation reduces total factor productivity, weakens fiscal sustainability, and complicates monetary policy. The recommended approach is a targeted reform rather than deregulation: regulatory sandboxes, SME exemptions, harmonised EU-level supervision, and mandatory review clauses.

Introduction

Artificial intelligence has become central to economic policy discussions over the past few years, and the reasons are not hard to see. The potential effects on productivity are large, the risks are real, and governments are under pressure to act. The European Union passed the AI Act in 2024, becoming the first jurisdiction to enact comprehensive AI legislation. In many ways, this was a serious achievement. The question this brief asks is not whether regulation was needed, but whether what was produced actually works.

Europe already trails the United States and China in AI investment by a significant margin. According to the OECD AI Policy Observatory, the EU attracted roughly a quarter of the venture capital that went into AI in the US in 2023, despite having a similar population and GDP. The Stanford AI Index for 2024 found that US firms produced around twice as many AI patents as the entire EU combined. These gaps have multiple causes, including differences in capital markets and university structures. But compliance costs play a role too, and the AI Act places a particularly heavy burden on smaller firms that do not have the legal or technical resources to absorb it.

This is not just a problem for the tech sector. From a macroeconomic perspective, regulation that persistently slows AI adoption has consequences for long-run productivity growth, fiscal sustainability, and how effectively monetary policy works. These effects are not immediate or dramatic. They build up slowly, which is probably why they tend to be overlooked in political debates focused on more visible short-term risks. The rest of this brief develops that argument and proposes a more proportionate regulatory approach.

Evidence: The Innovation Gap and Its Macroeconomic Implications

The Innovation Gap

Before moving to the theoretical framework, it is worth establishing what the data actually show. European AI adoption is weaker than the continent's economic size would suggest. A 2023 McKinsey report found that European firms adopt AI at notably lower rates than American ones across most industries. The number of AI unicorns based in Europe is a fraction of those in the US, and there is documented evidence of researcher and engineer migration toward Silicon Valley and, more recently, toward the Gulf states, where regulatory environments are considerably lighter.

The gap is also uneven within Europe itself. Germany and France, with larger economies and more established corporate research infrastructure, have managed to stay relatively competitive in certain AI sectors. The picture is worse for smaller member states, particularly in Central and Eastern Europe, where startup ecosystems are thinner and compliance costs take up a larger share of available resources. This matters because it suggests the AI Act may inadvertently widen existing development gaps within the EU, concentrating AI activity in a few larger economies while smaller ones struggle to keep up. It is worth noting that this kind of internal divergence is precisely what the single market was designed to reduce, so a regulation that reinforces it carries a real political cost on top of the economic one.

Supporters of the AI Act often point to what is sometimes called the Brussels effect: the idea that global firms will converge toward EU standards rather than maintain separate

compliance tracks, effectively giving Europe regulatory influence beyond its borders. There is some truth to this, and it has worked in areas like data privacy, where GDPR influenced regulatory approaches in other jurisdictions. The problem is that the AI Act is considerably more demanding than GDPR. While GDPR mainly required changes to data handling procedures, the AI Act mandates conformity assessments, technical documentation, and ongoing monitoring requirements that are far more resource-intensive. A large multinational can absorb this. A startup with fifteen engineers probably cannot. The Brussels effect argument may hold for established firms with existing compliance infrastructure, while at the same time creating barriers that prevent smaller and faster-moving firms from competing at all. That is a real tradeoff, and it tends to get glossed over in the policy debate.

A Macroeconomic Framework

The Solow growth model provides a useful starting point for thinking about the long-run implications. Output per capita is given by:

$$y = A \cdot f(k)$$

where y is output per worker, k is the capital-to-labour ratio, and A is total factor productivity, which captures how efficiently the economy uses its inputs. Once the economy reaches its steady state, further growth in output per capita can only come from increases in A . This is why technology matters so much in the long run.

AI fits into this framework as a positive shock to A . Technologies that make capital and labour more productive shift the production function upward and move the economy to a higher

steady-state output level. Regulation that delays or limits AI adoption pushes in the opposite direction. It keeps A below what it could have been, and because the effect compounds over time, even a relatively small drag on productivity growth translates into a meaningfully lower living standard over a generation. This is not a theoretical curiosity. Historical episodes, including the slow diffusion of electrification in the early twentieth century, suggest that the gap between a technology's invention and its broad economic impact often depends more on institutional and regulatory conditions than on the technology itself.

The fiscal consequences follow from this. The government's intertemporal budget constraint requires that the present value of future tax revenues is enough to cover future expenditures. In a two-period framework this is:

$$G_t + \frac{G_{t+1}}{1+r} = T_t + \frac{T_{t+1}}{1+r}$$

If long-run output is lower because productivity growth was suppressed, the tax base shrinks in future periods and this constraint becomes harder to satisfy. Governments have to either spend less or raise taxes. And since households in the representative agent model are forward-looking, they will anticipate these adjustments and reduce current consumption and investment accordingly. This is essentially the Ricardian equivalence logic running in reverse: expectations of worse future fiscal conditions depress economic activity before any actual tax increase takes place. The implication is that the fiscal cost of over-regulation is not simply a future problem. It can show up in lower investment today, through the expectations channel.

The connection to monetary policy is less direct but still relevant. The Taylor Rule describes how central banks set nominal interest rates:

$$i = \pi + r^* - b \cdot \frac{Y^* - Y}{Y^*} - c \cdot (\pi^* - \pi)$$

This rule depends on a reliable estimate of potential output Y^* . If AI regulation persistently suppresses productivity growth, estimates of Y^* will tend to be too low, and a central bank using those estimates may keep interest rates higher than the situation actually requires. This is not purely hypothetical. The ECB has faced significant criticism over its pace of adjustment during the post-2022 disinflation period, and the neutral rate r^* remains actively debated partly because structural changes, including digitalisation, make it hard to pin down.

None of this means the AI Act will trigger a recession or a fiscal crisis on its own. The magnitudes are uncertain and the effects would take years to show up clearly. But the direction is fairly consistent across all three channels, and the broader point is that regulatory design has macroeconomic consequences that go well beyond the firms being regulated.

Argument: Smarter Regulation Rather Than Less

To be clear from the start: the argument here is not that AI should be unregulated. There are genuine market failures involved. Algorithmic systems can reproduce and amplify existing biases in ways firms have little incentive to self-correct. AI applications in healthcare, credit scoring, and criminal justice raise serious accountability questions that markets will not resolve on their own. Some regulatory framework is both economically justified and necessary for public trust in AI systems used in high-stakes decisions.

The issue is with specific features of the current design. The compliance burden is not calibrated to firm size. A ten-person startup building a high-risk AI application faces essentially the same conformity assessment requirements as a large multinational with a dedicated legal team. That is disproportionate, and it has predictable consequences: smaller European firms are more likely to relocate, delay product development, or avoid high-risk application categories altogether. This is unfortunate precisely because those categories, including medical diagnostics, underserved credit markets, and public sector applications, are often where AI could do the most good. The regulation ends up being most restrictive in the areas where the potential social benefit is highest.

There is also a fragmentation problem. The AI Act is EU-level legislation, but enforcement is delegated to national authorities. The practical risk is divergent interpretations across member states, which creates compliance uncertainty for firms operating across borders. Since most firms of any scale operate in multiple EU countries, this effectively multiplies compliance costs rather than harmonising them. It works against the basic logic of the single market.

And there is no formal review mechanism. Technology moves fast in this space. A regulatory classification that was reasonable for the AI systems of 2024 may be poorly suited to those of 2029, but without a built-in review cycle there is no institutional process for updating it. Regulation tends to persist long after the conditions that justified it have changed.

Four Specific Recommendations

Based on the analysis above, this brief proposes four targeted reforms to the existing framework.

First, the EU should introduce regulatory sandboxes for startups and smaller firms. Companies below a defined threshold, perhaps under 250 employees or 50 million euros in annual revenue, should be allowed to develop and test AI applications under regulatory supervision without facing full conformity assessment from the outset. This model already exists in financial services. The UK Financial Conduct Authority's sandbox is generally seen as a workable way to balance innovation with oversight, and the principle transfers to AI reasonably well. The key point is that a sandbox is not a permanent exemption. It is a supervised testing environment that lets firms build a compliance record gradually, which is more proportionate than demanding full upfront compliance from organisations still in early development.

Second, compliance timelines should be tiered by firm size. Large firms have the resources to handle current requirements and should face them on the existing schedule. Smaller firms should have longer phase-in periods and access to simplified conformity tracks. This is not a special favour to the tech industry. It reflects a straightforward principle already applied in financial services and environmental regulation: the same absolute compliance cost means very different things depending on the size of the firm carrying it. Designing regulation as if all firms have the same capacity to comply is not neutral. It systematically advantages incumbents over new entrants.

Third, the EU should create a single supervisory authority for AI at the European level, along similar lines to the European Banking Authority. The current arrangement delegates enforcement to national bodies, which produces the fragmentation described above. Real harmonisation requires consistent interpretation, not just a common legal text implemented

differently in 27 countries. A centralised authority would also reduce legal uncertainty for firms trying to understand their obligations in ambiguous application areas, which is itself a barrier to investment.

Fourth, high-risk category classifications should face mandatory review every three years. The default assumption should be that regulators must justify the continuation of compliance requirements, not the other way around. This creates an institutional incentive to keep regulation in line with actual technological risk rather than letting it drift out of date. Given how rapidly AI capabilities are changing, a framework with no update mechanism seems likely to become increasingly misaligned with the risks it is supposed to address.

It is worth being explicit about what these recommendations do not propose. They do not call for dismantling the AI Act or removing safety requirements from genuinely high-risk applications. The argument is that regulation proportionate to actual risk and to the capacity of the regulated firms is more effective, and more politically sustainable, than a framework that imposes heavy costs on actors who cannot absorb them.

Conclusion

The EU AI Act is a serious attempt to govern a genuinely difficult technology. Its objectives are legitimate. The problem is that good intentions do not automatically produce good policy design, and some features of the current framework impose costs that are hard to justify against the benefits. From a macroeconomic perspective, persistently slowing AI adoption has

consequences for long-run productivity, fiscal sustainability, and monetary policy that build up gradually and tend to be underweighted in political debates focused on immediate risks.

Europe has some real advantages to build on. A well-regulated AI market could attract trust-sensitive applications in healthcare, finance, and public administration. But that only works if the underlying innovation ecosystem stays competitive, and that requires a regulatory framework that firms can actually operate within. The reforms proposed here are not radical. They keep the core of the AI Act in place while making it more proportionate, more consistent, and more capable of adapting over time. That seems like a more realistic path to achieving the Act's own goals than the current approach.

Note on Post-Presentation Feedback

During my presentation, one question was made about the monetary policy channel, and the truth is, it needs a more careful answer than the one I gave in the moment. My main text argues that if AI regulation blocks productivity growth, the estimates for potential output Y^* will be too low, which makes central banks keep the interest rates higher than necessary. The professor's question was: if the regulation really reduces what the economy can produce, shouldn't Y^* just go down to show this new reality? If that happens, the central bank is using all the information correctly, and there is no systematic error in the calibration.

Honestly, this is a very fair and well-grounded point. If the central bank's model updates in real time to show the slower growth from the rules, there is no normal measurement error for the output gap, the rates would be set right for the economy we actually have. Because of this, my argument in the brief needs to be explained with more precision.

The problem with monetary policy is better understood as a counterfactual issue, not just a standard measurement error. The central bank makes policy for the economy it observes. If the regulation stops AI adoption, both the actual output Y and the potential output Y^* are lower than they could be, and the estimated output gap might look small or almost zero. So, the rates look appropriate. But we are comparing against the wrong benchmark. The important question is not if the rates are correct for the suppressed economy, but if they are good compared to what we could produce without the regulatory drag. We never see that counterfactual Y^* in the data because that more productive economy never happened.

Also, there is a second mechanism that partially survives the professor's critique. To estimate potential output, we use models that look backwards, like the Hodrick-Prescott filter or structural models based on historical averages. When a big structural change happens, like everyone using a new general purpose technology, these models are slow to update because they mix recent data with a long historical baseline. During these transitions, even if the central bank uses all the information in good faith, it will underestimate the shifting productive frontier of the economy. This is not a failure of intention, but of the method, and we have seen it documented before with the IT boom in the 1990s and the secular stagnation debates after 2008. The same dynamic makes sense for AI, because the productivity effects are just starting and are super difficult to measure in real time.

To sum up: the professor is correct that a good central bank should track Y^* as it evolves, and my brief should not have implied a simple measurement error. The stronger version of the monetary argument is that suppressing AI with regulation creates a welfare cost that standard policy rules cannot detect, because those rules are calibrated against the realised economy instead of the counterfactual one. This distinction does not undermine my broader case for

smarter AI regulation, but it does require a more careful explanation of why monetary policy is affected.

References

- Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at work. *NBER Working Paper No. 31161*.
- European Commission. (2024). *Impact assessment accompanying the proposal for the AI Act*. SWD(2021) 84 final.
- IMF. (2024). *World economic outlook: Steady but slow: Resilience amid divergence*. International Monetary Fund.
- OECD. (2024). *OECD AI policy observatory: Trends and data on artificial intelligence*. OECD Publishing.
- Stanford HAI. (2024). *Artificial intelligence index report 2024*. Stanford University Human-Centered AI Institute.