Brian Williamson

Aligning regulation and AI

July 2024

COMMUNICATIONS CHAMBERS

About the Authors

Brian Williamson works at the intersection of technology, economics and policy. Clients have included governments, regulators, telcos and tech companies.

Disclaimer

This is an independent report funded by Google. The opinions offered herein are purely those of the authors. They do not necessarily represent the views of the client, nor do they represent a corporate opinion of Communications Chambers.

Contents

1. Executive summary 1
   - The prize 1
   - Aligning regulation and AI 1
   - Policy agenda to align regulation and AI 2

2. The prize – greater abundance and safety 3
   - General purpose technologies 3
   - The productive uses of AI will take time to emerge 4
   - AI is the invention of a method of invention 4
   - Productivity and jobs 4
   - AI may restore us to a goldilocks zone in terms of productivity growth 5
   - Limiting Baumol's cost disease 6
   - Safety benefits of AI 6

3. AI is subject to existing regulation, which will adapt 7
   - Applications of AI are regulated 7
   - Regulation/regulators will adapt 7
   - Case study – the UK domain specific approach to adaptation 7

4. What is it about AI that will drive the need for, and character of, adaptation? 9
   - Intelligence 9
   - Technology agnosticism 10
   - Scale 10
   - Solving for the equilibrium 11
   - Uncertainty and precaution 12

5. What is the best way to align regulation and AI? 14
   - Consider lessons from history 14
   - Start from baseline or counterfactual regulation 14
   - Focus on evolution and adaptation 14
   - Address constraints on AI use 14
   - Assess the risk of unintended regulatory consequences 15
   - Facilitate voluntary adoption of codes of practice and evaluation tools 15
   - Ensure AI input markets are responsive 16
   - Case study – the EU AI Act 17

6. Way forward globally 20

1. Executive summary

The prize

Artificial Intelligence (AI) is a general-purpose technology that is expected to contribute to productivity growth, the only sustainable source of per capita income growth. Higher productivity growth would help offset the headwinds of an ageing population. AI also has the potential to improve safety, and save lives, in areas ranging from healthcare to transportation.

AI also involves the invention of a method of invention in relation to science and technology. AI can therefore help achieve specific goals which require innovation, for example, contributing to advances in health care and facilitating decarbonization.

Aligning regulation and AI

Applications of AI are regulated by existing law and regulation. For example, if AI is utilised to identify a potential medicine, it will need to pass the same requirements in terms of efficacy and safety as existing drugs. It is therefore a false dichotomy to frame AI regulation as absent if not imposed via new law.

Law and regulation will need to evolve to align regulation and AI over time. It is not however realistic to attempt to anticipate all the changes that will be required and aim for a 'future proof' law of AI, it is also unnecessary.

We did not introduce a law of steam at the outset of the industrial revolution, and innovation without permission led to increases in the efficiency of steam engines and complementary innovations such as steam locomotives. Over time we introduced regulation – not of steam *per se*, but for boiler, rail and factory safety.

A challenge in relation to AI that we can anticipate is that existing law and regulation is often framed in terms of licensing of qualified individuals to assure quality and safety. We should adapt such regulation to ensure it is technology agnostic and that where AI is as safe or safer than the *status quo* approach, it can be adopted.

The UK is, for example, is adopting rules that permit autonomous boats alongside, and as an alternative to, existing rules regarding the need for qualified crew. Such reforms can be pursued case by case, but a comprehensive inventory of barriers to AI, and overarching enabling framework to remove them, may be justified.

We can also anticipate that AI, because of its general purpose nature, will not only be utilised by some to produce harms but will also be used to defend against them. Cybersecurity is an example where AI offers the prospect of stronger defense and potentially a reduction of harm overall.

Regulators and regulation constantly adapt to changes in technology, and much of this adaptation may be accommodated within existing legislative frameworks. Over time, issues may emerge that require legislative change.

Regulatory change may be driven by the fact that AI has predictive power. Returning to the medicines example, if AI is demonstrated to be capable of predicting efficacy and safety, then it may be possible to streamline medicines approval processes; shortening time to market and saving lives.

AI may also improve enforcement of some regulation (which may lead us to relax regulation somewhat if strict enforcement produces unintended harms); but may also increase some harms leading us to tighten or alter regulation.

The key is to adapt to align regulation and AI over time, not to legislate for AI *per se.* The following sets out a policy agenda for reform.

Policy agenda to align regulation and AI

Recognise existing regulation

AI applications are regulated under current laws. For example, where AI helps discover a new medicine that medicine will be subject to safety testing including clinical trials.

Consider net 'equilibrium' effects

AI will counter as well as facilitate harms e.g. in relation to cybersecurity. We should consider the net effect, seek to solve for the equilibrium and avoid regulation that could hand the advantage to attackers.

Adopt technology agnosticism to remove barriers to AI

Carry out an inventory of exiting law and regulation to identify provisions framed in terms of qualification requirements to perform tasks e.g. driver licensing and qualifications for the prescribing of medicines.

AI should be permitted to perform such tasks, provided it is as safe or safer than the *status quo*, subject to technology agnostic performance requirements. Reform to remove existing bias against AI could be case by case or potentially facilitated via a general enabling law.

Adapt over time to align regulation and AI

AI creates scope for more 'intelligent' market governance by providers, consumers and regulators. These opportunities should be permitted and seized. Some may require changes to law and regulation.

Technology agnosticism is part of the unlock here e.g. to enable consumers to utilise robot lawyers etc; but other opportunities will require reform beyond agnosticism e.g. AI-based prediction of molecular interactions may enable streamlined drug approval, but that would require a change to regulation.

Low-cost intelligence will leverage good and bad, creating scale effects that may justify regulatory change. Existing regulation may be calibrated around imperfect enforcement, if enforcement becomes cheap, it may be optimal to relax regulation. Existing regulation may also be calibrated around low levels of harm. If a harm is amplified, it may be optimal to tighten regulation - on a technology agnostic basis.

Disputes relating to liability and property rights will arise. The courts will rule on these. Some disputes may point to a need to amend existing legislative frameworks.

Don't regulate models or AI per se to permit open innovation

Recognising the value of innovation without permission, subject to technology agnostic downstream regulation, we should not regulate models or AI *per se*. We did not have a law of steam, but we did regulate rail safety; the approach to AI should be similar.

Industry will adopt standardised tests and benchmarks which should be voluntary and will evolve, though government may have a role in catalysing the development of open standards and benchmarks.

Ensure AI input markets are responsive to demand

We should ensure that the markets for inputs to AI; including energy, data and infrastructure/permitting are responsive to demand. This will support beneficial use of AI whilst limiting the impact of input scarcity, and price pressure, on the economy. AI itself may also facilitate market responsiveness.

2. The prize – greater abundance and safety

Artificial intelligence (AI) offers the prospect of greater abundance via higher productivity, socially beneficial applications and greater safety. That is the prize.

Productivity gains because users will only adopt AI where it offers prospective gains, and because AI is a general-purpose technology with a wide range of applications.

Socially beneficial applications because AI can help deliver on specific societal goals including improving health care and the green transition, both in terms of lowering costs but also opening new possibilities.

Safety gains because AI will have to meet or exceed existing standards, and in some contexts AI - as a complement or on its own – may greatly exceed performance under the *status quo*.

General purpose technologies

"Productivity isn't everything, but in the long run it is almost everything" Paul Krugman, 1992

Technical progress underpins productivity growth, and productivity growth is the only sustainable source of growth in income per capita and/or leisure.

Over the past two centuries productivity growth has increased income, leisure and contributed to longer life expectancy. Productivity growth, whilst involving the reallocation of labour, has not harmed overall employment.

General purpose technologies such as steam, electricity and connected computing underwent continuous improvement, saw widespread adoption, and made substantial contributions to productivity growth.[1] See figure 1.

Figure 1: Productivity growth contributions of previous general-purpose technologies

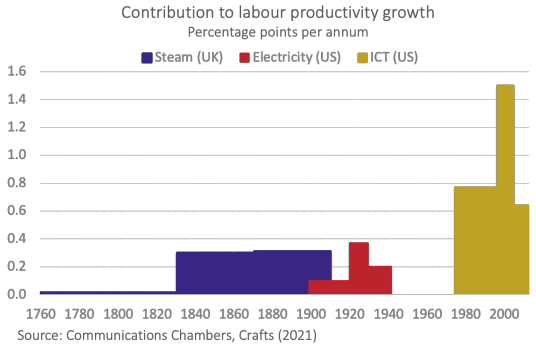

Source: Communications Chambers, Crafts (2021)

However, the process takes time as the technology has to reach a sufficient level of maturity to have widespread application (e.g. steam engines were initially inefficient), complementary innovations are required (e.g. steam locomotives) and because new forms of complementary industrial and social organisation take time to emerge.

The detailed implications of a general-purpose technology are hard to predict, even if we can partially predict the path of the technology itself (Moore's law, AI scaling laws). We do however know that general purpose technologies have in the past delivered substantial net benefits.

Consider electric power. At first it was used to do existing things, then factories were redesigned to utilise distributed motors. But it also made new goods and services possible including home appliances such as washing machines. In turn electric home appliances contributed to increased female labour force participation:[2]

[1] Crafts, September 2011, Artificial intelligence as a general-purpose technology: an historical perspective, Oxford Review of Economic Policy, Volume 37, Number 3. https://academic.oup.com/oxrep/article/37/3/521/6374675
[2] Jeremy Greenwood, 2019, How the appliance boom moved more women into the workforce. MIT Press. Extract: https://penntoday.upenn.edu/news/how-appliance-boom-moved-more-women-workforce

"In 1900, just 5 percent of married women worked; by 2000, this had risen to 61 percent. What can explain this jump? As female labor supply rose, time spent on housework declined. In 1900, the average household spent 58 hours a week on housework, including meal preparation, laundry, and cleaning—a figure that dropped to 18 hours in 1975."

We should show humility when thinking about new technology. The greatest benefits may be hard to foresee, whilst consumers, businesses and regulators can be expected to mitigate possible harms. In relation to potential harms, the new technology itself can be expected to help mitigate harms – it is an enabler. For these reasons, an optimistic stance is justified.

The productive uses of AI will take time to emerge

Whilst AI has great promise, widespread and productive use will take time to emerge because complementary innovations and changes to economic organisation will be required.

There are also institutional barriers to adoption including tasks and gatekeeper roles which only qualified individuals or approved institutions can perform. These frameworks should be reformed to accommodate the safe and productive use of AI.

AI is the invention of a method of invention

Not only is AI a general purpose technology, it is also the invention of a method of invention. AI has already been applied widely in science and technology, and can help achieve specific goals, for example, in bioscience and materials science.

Health and materials science applications of AI are relevant to two global challenges, an ageing population and the green transition.

In relation to health care AI can help improve access by reducing the costs of diagnostics and screening, manage overall costs through automation and help find new and more targeted medicines.[3]

The broader vision is that understanding biology, and therefore disease, disease prevention and treatment, is a challenge well suited to the application of AI:[4]

"I think biology can be thought of as an information processing system, albeit an extraordinarily complex and dynamic one." Demis Hassabis, 2021

In relation to health, AI is helping with the discovery of materials that are key to achieving affordable decarbonisation. For example, AI has helped identify candidate solid-state electrolytes for battery applications.[5]

Productivity and jobs

A naïve view is the productivity growth would mean less jobs. Yet technological progress opens new possibilities for work, and income growth driven by productivity growth stimulates demand for work throughout the economy.

Historically the net impact of labour productivity growth on employment has been negligible, as shown in Figure 2 for several Western European countries over the period 1870-1998.

[3] The Economist, 30 March 20924, Health and AI - A new prescription. https://www.economist.com/technology-quarterly/2024-03-30

[4] Demis Hassabis, November 2021, Introducing Isomorphic Labs. https://www.isomorphiclabs.com/articles/introducing-isomorphic-labs

[5] Chen *et al*, January 2024, Accelerating computational materials discovery with artificial intelligence and cloud high-performance computing: from large-scale screening to experimental validation. https://arxiv.org/abs/2401.04070

Figure 2: Productivity growth and employment

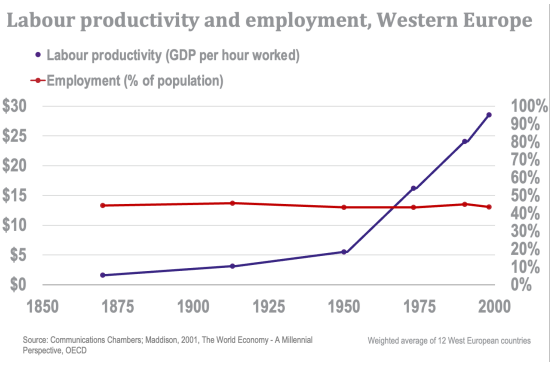

However, work changed radically over the period with new categories created and some categories almost eliminated. Hours worked also declined, almost halving over the period i.e. half of the 18-fold gain in productivity per hour worked translated into income with the rest going to leisure.

Considering the labour share of national income, Crafts (September 2021) found that during the period 1770 to 1860 it was essentially unchanged, rising from 61% to 62%, but that it showed a very short and modest decline at the begging of the nineteenth century.

More recently the labour share has declined in the US and globally.[6] Views differ as to what impact AI might have on wages, the income share and the income distribution.[7] [8]

Adverse distributional consequences are a possibility, and some individuals will be worse off (almost any change makes some worse off, even if it grows the overall pie). However, income and other forms of social support are more developed than during early waves of general purpose technologies.

AI may restore us to a goldilocks zone in terms of productivity growth

One way to think about this is to consider the pace of disruption, of creation and destruction of jobs and firms, versus historical rates. The rate of creation and destruction of jobs and firms has declined in OECD countries[9], but is nevertheless substantial.

For the US over the period 1980 to 2009 the annual rate of job creation and destruction was around 15%[10] (which would equate to roughly 500 million jobs created and destroyed per annum globally). When estimates of the potential for automation from AI are discussed the high background rate of job creation and destruction should be kept in mind.

Another consideration is that a counterfactual of continued low productivity growth, which many countries have experienced over the past two decades, is hardly attractive.

An acceleration of productivity growth would not only help address 'the cost of living crisis' (the slowdown recast in terms of prices rather than real incomes) but would also help with the challenges of high national indebtedness, an ageing population and declining workforce coupled with the transition to reduced use of fossil fuels. A revival of economic and social progress due to the application of AI may also help mitigate counterproductive zero-sum thinking, a symptom of low growth.[11]

Rather than thinking of AI as ushering in an era of radical and unwelcome disruption, we should consider the potential for AI to restore productivity growth to the goldilocks zone with

[6] Karabarbounis, Spring 2024, Perspectives on the Labour Share. Journal of Economic Perspectives, 38(2). https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.38.2.107
[7] Autur, February 2024, Applying AI to Rebuild Middle Class Jobs. https://www.nber.org/papers/w32140
[8] Acemoglu & Johnson, May 2024, Learning from Ricardo and Thompson: Machinery and Labor in the Early Industrial Revolution, and in the Age of AI. https://www.nber.org/papers/w32416
[9] https://www.oecd-ilibrary.org/science-and-technology/declining-business-dynamism_77b92072-en
[10] https://www.journals.uchicago.edu/doi/full/10.1086/663154
[11] Sahil Chinoy *et al*, September 2023, Zero-Sum Thinking and the Roots of U.S. Political Divides. https://www.nber.org/papers/w31688#fromrss

which society has coped well in the past, and successive generations have benefited.

Limiting Baumol's cost disease

The relative cost of things that are hard to automate will rise as productivity growth drives up real wages and as AI competes for inputs with other activities. To counteract this effect barriers to AI adoption in areas such as education and health care, which have historically seen relative cost growth - Baumol's cost disease[12] - should where feasible and consistent with safety be reduced. Further, we should ensure that markets for inputs to AI are responsive.

Safety benefits of AI

AI will have to meet or exceed existing safety requirements and has the potential to significantly exceed existing human standards of performance.

AI augmented or autonomous vehicles should ultimately be safer than those driven by humans, whilst the application of AI to health care has the potential to improve access (contributing to healthier outcomes for all) and safety via AI augmented decision making.

AI is not risk-free, but neither are we, and our performance is biased, noisy, difficult to align with institutional and societal goals and corruptible.

In conclusion, AI has the potential to contribute to productivity growth, socially beneficial progress in areas such as health and the green transition and improvements in safety. Rather than fearing unprecedented rates of disruption due to AI, we should focus on the opportunity for AI to help address pressing challenges and offsets the headwinds of an ageing population.

[12] Timothy B. Lee, May 2017, William Baumol, whose famous economic theory explains the modern world, has died. Vox. *https://www.vox.com/new-money/2017/5/4/15547364/baumol-cost-disease-explained*

3. AI is subject to existing regulation, which will adapt

Applications of AI are regulated

If AI is used to identify a prospective medicine, then that prospective medicine will be subject to existing appraisal and safety requirements including clinical trials.

Medicines are not the only example, there are myriad regulations covering other vertical segments of the economy as well as horizontal laws which govern all activity (e.g. consumer protection, data protection, non-discrimination etc). These laws and regulations will apply to goods and services utilising AI.

Another illustrative example involves a democratic consultant indited on felony charges of voter suppression and impersonation after commissioning deepfake robocalls impersonating President Joe Biden's voice (on behalf of a democratic candidate running against Biden). The FCC have also proposed a $6 million fine for violation of the Truth in Caller ID Act.[13]

It may be that amended or new law is required to address potential issues in relation to the misuse of AI. However, existing law may adequately cover most instances of potential misuse.

Regulation/regulators will adapt

Regulation is routinely adapted in response to technological change. Technological change may also trigger disputes over property rights and/or liability issues, which in the first instance are addressed by the courts; but which could ultimately point to the need for legislative or regulatory change.

To the extent that existing regulation is framed in a technology agnostic manner or can readily be re-interpreted (for example via the decisions of an independent regulator) this process is in principle straightforward.

Given that AI is a general purpose technology which has the potential to offer large benefits, a more deliberative review (rather than simply wait and see and adapt) of existing regulation may be justified. There may also be common issues that need to be addressed, and which may benefit from co-ordination amongst regulators and potentially guidance.

In instances where existing regulation cannot adequately accommodate AI based services legislative change may be required. The need for specific legislative change can be expected to emerge from detailed review and dialogue across different sectors and in relation to existing horizontal law i.e. it is not something well suited to a generic 'law of AI'.

Likewise, a 'law of steam' would likely have proved counterproductive at the outset of the industrial revolution and would have failed to address real issues that arose including factory and rail safety. Important applications often significantly lag the underlying technology – the first public stream railway was in 1825, 56 years after James Watt's steam engine patent.

Case study – the UK domain specific approach to adaptation

The UK government white paper "A pro-innovation approach to AI regulation" set out an approach along these lines:[14]

[13] Arstechnica, May 2024, Democratic consultant indicted for Biden deepfake that told people not to vote. *https://arstechnica.com/tech-policy/2024/05/democratic-consultant-indicted-for-biden-deepfake-that-told-people-not-to-vote/*

[14] Department for Science, Innovation and Technology, August 2023, A pro-innovation approach to AI regulation. *https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper*

"New rigid and onerous legislative requirements on businesses could hold back AI innovation and reduce our ability to respond quickly and in a proportionate way to future technological advances. Instead, the principles will be issued on a non-statutory basis and implemented by existing regulators. This approach makes use of regulators' domain-specific expertise to tailor the implementation of the principles to the specific context in which AI is used."

Following the White Paper individual regulators including the Competition and Markets Authority, the Financial Conduct Authority, the Information Commissioners Office and Ofcom have considered how AI may impact their approach to regulation.

Further, the Digital Regulation Cooperation Forum (DRCF), with the above regulators as members, launched the AI and Digital Hub ("the Hub") pilot to support innovators working on AI or digital products in navigating existing regulation.[15]

An illustration of how some issues may need to be addressed by targeted legislative change is provided by autonomous ships.[16] Whilst autonomous ships already exist in the UK, these ships' operations have so far been addressed through exemptions and equivalences to the existing regulations for non-autonomous ships.

However, the exemption is a temporary measure, and the government proposes amendments to current primary legislation framework to provide the powers to regulate autonomous ships and a predictable framework for the growth of the industry.[17]

[15] Digital Regulation Cooperation Forum (DRCF), April 2024, DRCF AI and Digital Hub launched to support innovation and enable economic growth. *https://www.drcf.org.uk/news-and-events/news/drcf-ai-and-digital-hub-launched-to-support-innovation-and-enable-economic-growth2*

[16] BBC, 7 March 2024, Robot ships: Huge remote controlled vessels are setting sail. *https://www.bbc.co.uk/news/science-environment-68486462*

[17] Department for Transport, September 2023, Maritime Autonomy and Remote Operations – Impact Assessment. *https://assets.publishing.service.gov.uk/media/650198b239d9f100143d3a41/fot-regulatory-review-maritime-autonomy-and-remote-operations-final-impact-assessement.pdf*

4. What is it about AI that will drive the need for, and character of, adaptation?

The following considerations will drive the need for, and character of, adaptation required to align regulation and AI:

- *Intelligence.* AI increases the availability of the kinds of intelligence AI is good at. In a sense everything about AI, and the need for alignment of regulation and AI, flows from this.
- *Technology agnosticism*. Existing law and regulation may be framed in ways that do not permit beneficial applications of AI because the possibility of artificial intelligence was not anticipated when the rules were made.
- *Scale*. AI will lower the costs of activities requiring the sorts of intelligence AI offers. This will enable activities at scale that were previously economically infeasible.
- *Solving for the equilibrium*. In response to the first-round impact of AI people and markets will adapt. Before adapting regulation, one should either attempt to 'solve for the equilibrium'[18], or wait and see, since problems may resolve themselves, be preferred to feasible but imperfect intervention or change in character in ways that change the appropriate intervention.
- *Uncertainty and precaution*. We don't know what the most important applications will be, what new harms may arise and persist or what harms AI may help address. The notion of precaution is ambiguous in principle. Should we be more cautious about the risk we forego benefits including safety benefits, or about avoiding avoidable risks?

Intelligence

AI is increasing the availability of the kinds of intelligence AI is good at. In a sense everything about AI, and the need for alignment of regulation and AI, flows from this, but it needs unpacking to offer useful guidance.

AI will allow some existing tasks to be performed at lower cost and/or more safely, potentially by a large margin versus the *status quo*.

Consider autonomous vehicles. The prospect of low accident rates, and improved access for those unable to drive, are unambiguously good – but they will only be realised if we adapt regulation to permit autonomous vehicles i.e. technology agnosticism with respect to natural versus artificial intelligence 'behind the wheel'.

Consider low-cost image recognition. Our view arguably depends on the context, but the underlying effect is one of scale. It is also helpful to note that image recognition may be utilised by platforms (or consumers) to impose community (or personal or family standards), by platforms to comply with regulation or by regulators to enforce regulation, for example, to identify advertising that may not comply with advertising standards[19].

Everyone can utilise AI, and regulators should adapt regulation with the possibility that responsibility for 'enforcement', for ensuring that markets work well, could shift versus a world in which certain types of intelligence were costly and scarce.

AI will also allow new things, in particular prediction and complex information processing

[18] A common prompt on the blog Marginal Revolution, framed another way '…and then what?'. *https://marginalrevolution.com*

[19] ASA, 2023, AI-assisted collective ad regulation. *https://www.asa.org.uk/resource/ai-assisted-collective-ad-regulation.html*

that were either not possible or too costly previously.

For example, our current approach to testing the safety of new medicines is, in part, based on successive testing via animal trials and clinical trials. We are testing for efficacy and adverse effects, ultimately across a large and diverse group of people, because we cannot predict outcomes with any precision.

AI is helping with drug candidate discovery, with candidate drugs following the standard approval procedure. But given advances in AI-based prediction for complex biological interactions (e.g. AlphaFold 3[20]) there may be scope in future to shorten the drug approval process, once such predictions have been validated.[21] This would require regulatory adaptation.

AI may also support individual consumers, for example, reducing information asymmetries by comparing complex offers and ratings to make a reasoned recommendation; and reducing barriers to access to legal support and dispute resolution. Examples of the latter include DoNotPay![22] and Contend[23]. Such 'robot lawyer' services depend on a permissive regulatory environment. Generative AI apps may also act as a consumer assistant. Market dynamics, and the need for and nature of consumer protection regulation, may adapt as a result.

Technology agnosticism

Technology agnosticism is desirable because it leaves open scope to implement the most efficient and/or safe combination of natural and artificial intelligence.

AI starts with a disadvantage since some tasks are reserved for people with specific qualifications via professional and occupational licensing, for example only certain qualified individuals can prescribe a medicine or crew a ship. It is estimated that nearly 30 per cent of workers in the US[24] need a license to perform their job, and 22 per cent of those in the EU[25]

We should consider opening such tasks up to AI, provided outcomes as good or better can be assured. Reform to reduce the scope of occupational licensing, which has outgrown the public interest goal of public safety[26], would also achieve the same goal and could be pursued in parallel.

Delivering technology agnosticism will ultimately require detailed re-examination of existing rules. Temporary exemptions might be granted to trial new approaches, or an across the board enabling act or guidance issued to facilitate technology agnosticism could be pursued.

On the other hand, efforts to legislate for AI *per se* may exacerbate the problem of non-neutrality by placing new requirements on artificial but not natural intelligence based systems and services.

Scale

AI lowers the cost of activities dependent on the sorts of 'intelligence' AI offers and will enable

[20] Google DeepMind, 8 May 2024, AlphaFold 3 predicts the structure and interactions of all of life's molecules. *https://blog.google/technology/ai/google-deepmind-isomorphic-alphafold-3-ai-model/*
[21] Bloomberg Television, 9 May 2024, Google DeepMind CEO on Drug Discovery, Hype, Isomorphic. At 11:36. *https://www.youtube.com/watch?app=desktop&v=RIrnMVDM_N8*
[22] A16Z Podcast, March 2023, The Robot Lawyer Resistance. *https://a16z.com/podcast/the-robot-lawyer-resistance/*
[23] *https://contend.legal*
[24] Hershbein, Boddy, and Kearney, January 2015. Nearly 30 Percent of Workers in the U.S. Need a License to Perform Their Job: It Is Time to Examine Occupational Licensing Practices. *https://www.brookings.edu/articles/nearly-30-percent-of-workers-in-the-u-s-need-a-license-to-perform-their-job-it-is-time-to-examine-occupational-licensing-practices/*
[25] Koumenta and Pagliero, 2018, Occupational licensing in the European Union, Coverage and Wage Effects. CEPR Discussion Paper. *https://cepr.org/publications/dp12577*
[26] The Economist, February 2018, Occupational licensing blunts competition and boosts inequality. *https://www.economist.com/united-states/2018/02/17/occupational-licensing-blunts-competition-and-boosts-inequality*

activities at scale that were not previously feasible.

In some instances, something that was tolerated or ignored at small scale may become a problem that we decide to address with new or altered law. For example, fake images or the potential for low-cost surveillance at scale. Ideally such law would be framed in a manner that was AI agnostic, but it may be motivated by AI.

In other instances, enforcement of existing rules by regulators will become much easier. This is helpful but may point to a need to reform the underlying rules e.g. if the potential to enforce copyright at scale discouraged creative use of content (copyright law may not have been designed with enforcement at scale in mind and it may be unclear what in a large data set is copyrighted[27]).

Related to scale is access. Some services that were previously rationed or constrained by affordability may become much more widely available and/or affordable with AI. Professional services such as health care and legal support are examples.

Solving for the equilibrium

The discussion above has focussed on developments we might reasonably predict, and which illustrate the need to adapt to align AI and regulation.

However, at least implicit in the discussion up to this point is that AI is tool that can be utilised by all, that some initial impacts may be large, and that the final 'equilibrium' will depend on the interaction of multiple agents utilising AI.

The shift in thinking required from first round impacts to equilibrium thinking isn't just second order, it may change the anticipated direction, or nature, of impact of AI.

We should therefore either aim to 'solve for the equilibrium' in advance; or wait and see. And if one thinks AI, more intelligence, is going to produce net harm; be clear about why one thinks more intelligence is bad when we generally tend to think more intelligence is good.

AI is a tool, and it will be used offensively and defensively. For example, software facilitated email spam, but various countermeasures, now including AI, enable Gmail to block more than 99.9% of spam, phishing attempts and malware.[28]

Cybersecurity is another example, with AI identified as a tool that could amplify harm:[29]

"AI's impact is not only amplifying old digital problems and creating new societal challenges… The advancement of AI technologies is introducing significant security risks by increasing vulnerability to sophisticated cyber-attacks."

However, this assessment neglects consideration of how AI might be utilised for cybersecurity defence, and whether AI would on balance therefore increase or decrease cybersecurity risks. Assessments that have considered the potential for AI enhanced defence paint a different picture:[30]

"We believe AI affords the best opportunity to upend the Defender's Dilemma, and tilt the

[27] Coalition for Creativity, September 2023, C4C's Perspective on the EU AI Act: Copyright in Real Life is Messy and AI Discussions Are Not Helping. *https://coalition4creativity.org/2023/09/28/c4cs-perspective-on-the-eu-ai-act-copyright-in-real-life-is-messy-and-ai-discussions-are-not-helping/*

[28] Google, October 2023, New Gmail protections for a safer, less spammy inbox. *https://blog.google/products/gmail/gmail-security-authentication-spam-protection/*

[29] Floridi, April 2024, The Ethics of Artificial Intelligence: exacerbated problems, renewed problems, unprecedented problems. *https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4801799*

[30] Google, February 2024, Secure, Empower, Advance - How AI Can Reverse the Defender's Dilemma. *https://services.google.com/fh/files/misc/how-ai-can-reverse-defenders-dilemma.pdf*

scales of cyberspace to give defenders a decisive advantage over attackers."

Other examples where thinking about first round impacts alone may distract from the ultimate policy challenge include disinformation. As The Economist noted:[31]

"Encouragingly, technology is as much a force for good as it is for evil. Although ai makes the production of disinformation much cheaper, it can also help with tracking and detection. Even as campaigns become more sophisticated, with each spreader account varying its language just enough to be plausible, ai models can detect narratives that seem similar. Other tools can spot dodgy videos by identifying faked audio, or by looking for signs of real heartbeats, as revealed by subtle variations in the skin colour of people's foreheads."

Further, commenting on nation states attempts to interfere in elections, Mark Zuckerberg noted:[32]

"They're sophisticated. They have a lot of technology. In those cases, I still think about the ability to have our AI systems grow in sophistication at a faster rate than theirs do. It's an arms race but I think we're at least winning that arms race currently."

Further, if misinformation did nevertheless increase then information, including misinformation, would be devalued, not necessarily believed.[33] That may still be a problem, but it is a different problem, and one markets may help address since the value of differentiating yourself by promoting a brand that utilises and verifies trusted sources is higher in an environment where disinformation is cheaper.

Before adapting regulation to anticipated harms we should think through the potential equilibrium when 'intelligence' is utilised by all, when business models have had an opportunity to adapt. And if that is too hard, there may be an argument for waiting to see.

Uncertainty and precaution

The development of AI is itself uncertain, though scaling laws offer some insight. However, even if knew how a general purpose technology was going to develop, its implications will depend on complementary innovations (e.g. rail for steam) which are largely unknown.

The implications will also depend on the interaction between agents utilising AI, and as the previous section on disequilibrium illustrated, this may change not just the magnitude, but the direction of impact.

A further indication of uncertainty is that the focus in relation to AI has shifted over time. For example, when the EU AI Act was initially proposed generative AI was not a focus.

We are dealing with radical uncertainty, with unknown unknowns. So, what should we do?

We should anticipate that we will need to adapt regulation, whilst we should also recognise that any hope of a future-proof 'law of AI' is simply hubris.

Had we legislated for a law of steam at the outset of the industrial revolution it may have impeded or delayed beneficial applications of steam, whilst failing to address the specific safety issues that emerged over time.

An adaptive approach is more likely to be effective, but should we nevertheless exercise caution before proceeding with new

[31] The Economist, 2 May 2024, How disinformation works—and how to counter it. https://www.economist.com/leaders/2024/05/02/how-disinformation-works-and-how-to-counter-it
[32] Dwarkesh podcast, 18 April 2024, Mark Zuckerberg. https://www.dwarkeshpatel.com/p/mark-zuckerberg
[33] Joshua Gans, November 2023, AI and Election Misinformation – what should we be worried about? https://joshuagans.substack.com/p/ai-and-election-misinformation

applications? After all, this is the stance we adopt in relation to a new medicine.

However, the correct stance depends on the expected probability distribution of benefits and costs, and on how we learn over time.

When it comes to new medicines, we run randomised controlled trials, which take time, to check for unintended side-effects. This involves trade-offs, and when we know for certain that people would otherwise die new treatments may be permitted with less assurance of safety.

In relation to AI, Gans (2024) considers the question of whether such a process would work, and concludes that it depends:[34]

"For any AI directed at predicting physical things, the answer is probably yes, but these are also areas where there exist product safety regulations."

"But what about unintended consequences, bad actors or equilibrium consequences? Experiments in the lab or field are good at providing first-order impacts, but what about impacts that only reveal themselves as uncontrolled adoption takes place? Suppose you look at the concerns regarding AI. They are almost all in this latter category: A bad actor takes an LLM and uses it to engage in cyberattacks, or the continued use of algorithms leads to the undermining of democracy.

The point here is that there are real reasons to believe that only true "market" tests can reveal these issues with AI."

Precaution can imply a 'bias' towards acceleration:[35]

"When such learning is by doing -- based on the real-world adoption of AI -- this generally favours acceleration of AI adoption to surface and react to potential harms more quickly."

We also expect that the application of AI will produce widespread economic and safety benefits, including lives saved – perhaps a great many - in areas including medicine and transportation where human error is a key contributor to risk.

Some interpretations of precaution may also preclude or hamper the development of open-source models, thereby foregoing the potential innovation, safety and security benefits from allowing open-source AI.[36]

Delaying AI may therefore be costly and reveal little regarding risks that are not already managed by existing regulation.

Further, unlike medicines the risks associated with AI may materialise whether a given country delays the application of AI, indeed overall risk may be increased if AI itself could have been utilised defensively and this is precluded or delayed.

[34] Joshua Gans, September 2023, Should we slow down the development and adoption of AI? https://joshuagans.substack.com/p/should-we-slow-down-the-development

[35] Joshua Gans, January 2024, How learning about harms impacts the optimal rate of artificial intelligence adoption. https://www.nber.org/papers/w32105

[36] Eiras *et al*, May 2024, Risks and Opportunities of Open-Source Generative AI. https://arxiv.org/pdf/2405.08597

5. What is the best way to align regulation and AI?

Consider lessons from history

We did not legislate to regulate other general purpose technologies including steam, electricity and computing *per se.*

Existing rules applied to successive technologies (rail safety applies to electric as well as steam trains), and new rules were developed where needs arose ranging from boiler safety to electrical wiring standards and data protection.

An approach of learning by doing, of adapting existing regulation and potentially legislation as issues emerge, was followed; and has merit given the level of uncertainty regarding the evolution and diffusion of general-purpose technologies.

History also points to the tendency for society to initially act with a bias against the new, to attribute harms to new things where they are either imagined (e.g. 'bicycle face'[37]) or the linkage is so indirect that it would be wiser to focus on the harm itself. Douglas Adams (1999) put the latter as follows in relation to the internet:

"Newsreaders still feel it is worth a special and rather worrying mention if, for instance, a crime was planned by people 'over the Internet.' They don't bother to mention when criminals use the telephone or the M4, or discuss their dastardly plans 'over a cup of tea,' though each of these was new and controversial in their day."

We should recognise our bias, step back and assess whether initially perceived problems stand up to scrutiny.

Start from baseline or counterfactual regulation

The baseline or counterfactual is that applications of AI are regulated by specific vertical law and regulation (e.g. medicines safety) and horizontal law (e.g. non-discrimination requirements).

Applications of AI are also 'regulated' by consumer preferences and choices and by market institutions, for example, mechanisms implemented by service providers to promote trust in their service. Consumers and businesses will also utilise AI to help them make good choices, and promote trust and achieve regulatory compliance, respectively.

Yet much that has been written about AI ethics, alignment and regulation does not start with this baseline, rather it assumes a blank slate. This is neither realistic nor informative.

Focus on evolution and adaptation

Both delegated 'independent' regulation and consumer choices and market institutions will adapt and evolve with AI. To an extent alignment of regulation and AI will occur endogenously.

Endogenous evolution has the advantage that it allows continuous adaption as AI develops and is applied in different applications. However, it has limits.

Address constraints on AI use

Endogenous adaption has limits, it is constrained by existing legislation and new approaches may be required. Whilst there is considerable uncertainty, and a wait and see approach is justified to an extent whilst we discover the

[37] Joseph Stromberg, March 2015, "Bicycle face": a 19th-century health problem made up to scare women away from biking. Vox. https://www.vox.com/2014/7/8/5880931/the-19th-century-health-scare-that-told-women-to-worry-about-bicycle

limited of endogenous adaption, some issues that will require attention can be identified.

Existing law and regulation are framed with humans (including their motivations, inscrutability, behavioural biases and cognitive constraints) and human institutions, which in part are designed to counteract our individual failings, in mind.

First, existing law and regulation may prevent the productive and safety enhancing application of AI. This is the case where only qualified humans and designated institutions can perform certain tasks e.g. prescribing a medicine or crewing a boat. It may be desirable to introduce safety equivalent outcome-based requirements which would permit AI. This can be anticipated and might be best tackled via comprehensive reform.

Second, the detailed ways in which we verify compliance, for example, with non-discrimination laws, might not be the best possible fit with AI. For example, with AI simulations are possible to test compliance, training is more transparent, and learning can be directly shared. It may also be the case that rules that may be helpful for humans e.g. hiding the ethnicity or gender of job applications to reduce scope for discrimination; would be unhelpful with AI if the AI model could use such information to reduce bias.[38] These are questions of detail and might be best left to sector specific regulators in the first instance.

Third, regulatory capture may have occurred or may be a risk either by non-AI based incumbents or AI providers. There is no single solution to this problem, but at times governments change governing legislation and regulatory institutions to overcome such problems, where technological change makes innovation in some area a priority. An example is the introduction of frameworks to permit app-based ride hailing.

Assess the risk of unintended regulatory consequences

Regulating early, as a technology and its applications are evolving rapidly, increases the risk of unintended harmful consequences. Not only may innovation which could enhance productivity be slowed as a result, but safety benefits may also be delayed. Good practice can reduce the scope for unintended consequences. In particular:

- Consider counterfactual regulation before adding additional rules, is it adequate, can it adapt, can regulators utilise AI to do so?
- Consider the evolution of the market counterfactual, how will markets adapt, to what extent might market participants address emerging challenges, can consumers and firms utilise AI to do so?
- Are requirements for AI proposed that humans themselves could not comply with, is the trade-off with technology agnosticism justified and is there a risk that something valuable will be foregone as a result?

Several potential unintended consequences are considered in relation to the EU AI Act are considered later in this section.

Facilitate voluntary adoption of codes of practice and evaluation tools

In a dynamic environment the risk of unintended consequences can be reduced by working with market participants on measures including codes of practice, voluntary standards and evaluation tools. These can be adopted quickly, adapted as circumstances change and leave open the possibility of competing ideas.

An example is The White House voluntary AI Commitments entered into by several leading

[38] Kleinberg and Mullainathan, August 2018, Simplicity Creates Inequity: Implications for Fairness, Stereotypes, and Interpretability. *https://arxiv.org/pdf/1809.04578.pdf*

AI companies in July 2023.[39] Another is commitments entered into at the UK AI Safety Summit in November 2023.[40]

In relation to model evaluation and benchmarking, there are many tools in existence, and they will continue to evolve as model capabilities and our understanding develop.

An example of an initiative designed to encourage open evaluation was the release by the UK AI Safety Institute of a testing platform "Inspect" to strengthen AI safety evaluations.[41] The software was released under an open source licence, which not only means it is freely available but that others can help adapt and build upon the platform.

Ensure AI input markets are responsive

AI utilises computing, electricity, data and human capital as key inputs. To maximise the benefits from AI it is important that the markets for these inputs are responsive to demand, up to the point where the marginal benefit of additional AI 'intelligence' equates to the marginal cost of the input including any externality.

For electricity both generation and grid connections may prove a constraint on AI. The remedy includes consideration of permitting requirements for generation and grid, technology agnosticism regarding decarbonisation and ideally carbon pricing or permits rather than constraints on generation. Long-term commercial relationships between providers of clean energy and compute infrastructure may help drive an increase in abundance of electricity consistent with decarbonisation by guaranteeing demand "*for developers of promising technologies like long-duration energy storage, clean hydrogen, next-generation geothermal and nuclear energy*".[42]

For data, the challenge is in recognising that the trade-offs have changed, and that data protection should be focussed on protection rather than limitation. Demirer, Hernández and Peng (February 2024) found that the GDPR had reduced data storage and processing and raised costs:[43]

"EU firms decreased data storage by 26% and data processing by 15% relative to comparable US firms, becoming less "data-intensive." To estimate the costs of the GDPR for firms, we propose and estimate a production function where data and computation serve as inputs to the production of "information." We find that data and computation are strong complements in production and that firm responses are consistent with the GDPR, representing a 20% increase in the cost of data on average."

The availability of data for training AI models may also be constrained by copyright requirements. In considering what constraints should apply the costs of overprotection of

[39] The White House, July 2023, Ensuring Safe, Secure, and Trustworthy AI. *https://www.whitehouse.gov/wp-content/uploads/2023/07/Ensuring-Safe-Secure-and-Trustworthy-AI.pdf*
[40] November 2023, Chair's Summary of the AI Safety Summit 2023, Bletchley Park. *https://www.gov.uk/government/publications/ai-safety-summit-2023-chairs-statement-2-november/chairs-summary-of-the-ai-safety-summit-2023-bletchley-park*
[41] AI Safety institute, May 2024, AI Safety Institute releases new AI safety evaluations platform. *https://www.gov.uk/government/news/ai-safety-institute-releases-new-ai-safety-evaluations-platform*
[42] The Economist, 5 May 2024, Big tech's great AI power grab. *https://www.economist.com/business/2024/05/05/big-techs-great-ai-power-grab*
[43] Demirer, Hernández and Peng, February 2024, Data, Privacy laws and firm production: evidence from the GDPR. *https://www.nber.org/papers/w32146*

copyright, which have arguably risen with generative AI, should be considered.[44]

Case study – the EU AI Act

The EU specific challenge

It is helpful to consider the context surrounding the introduction EU AI Act, since this is relevant to other nations that might contemplate mimicking aspects of the EU AI Act.

The Act is designed to address an EU specific challenge within the constraints of EU institutions. The challenge is to head off the risk of disparate member state legislation introducing barriers to scaling applications EU-wide, an issue that is not necessarily relevant to other nations.

However, the EU AI Act alone will not deliver a single regulatory environment for AI since applications will still be subject to national law and regulation at the application specific level, and these laws and regulations are neither necessarily harmonised EU-wide, nor adapted to reconcile regulation and AI. Those challenges remain, and the EU AI Act is silent on them.

In relation to EU specific constraints, the EU has limited scope to do anything other than regulate since it has no powers of taxation, a budget of only around 1% of EU GDP and limited scope to pursue a bottom-up adaptive approach to reconciling regulation and AI since many of the relevant 'competences' reside at member state level.

The EU legislative process also has its own specifics in terms of the role of parliament, the European Commission and member states represented via the Council of Ministers. It is also motivated by a desire to legislate to enshrine European values. In the case of the EU AI Act this has involved compromises and resulted in a mix of general and rather specific clauses, some of which are unlikely to stand the test of time.

For the EU, the choice was between a horizontal law of AI or having little, if any, influence over how the market for AI applications developed in the EU. Other nations have other options and should consider them.

Legislating in a dynamic environment

The EU AI Act incorporates some flexibility, some scope for interpretation, in its application. The aim was to be 'future proof'.

But as the notion of a future proof law of steam at the outset of the industrial revolution illustrates, aiming for future proof law of AI is a Sisyphean task.

During drafting, the EU AI Act was modified to accommodate technological change. The emergence of generative AI led to the adddition of a regime focussed on general purpose AI models which is separate from the rest the regime which focusses on AI systems (essentially applications). This introduces a risk of overlapping model and systems regulation, which should be guarded against.

The introduction of general-purpose AI models was accompanied by a provision for a training compute threshold (the US 2023 AI Executive Order also includes a compute threshold)[45]. Whilst intended as a proxy for capability, a compute threshold is an odd basis for regulation.

Capabilities are not obviously correlated with risk, and compute is likely to prove a poor proxy for capability (more capable models are being compressed and techniques developed to allow

[44] Bertin Martens, April 2024, Economic arguments in favour of reducing copyright protection for generative AI inputs and outputs, Working Paper Issue 09/2024. *https://www.bruegel.org/sites/default/files/2024-04/WP%2009%20040424%20Copyright%20final_0.pdf*

[45] The White House, October 2023, Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. *https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/?ref=blog.heim.xyz*

them to run on low power devices). More capable models may be more accurate, easier to align with specific objectives and constraints and more capable in defensive roles.

Compute thresholds for regulation are not future proof, and appear likely to have perverse effects, potentially increasing overall risk.

Another aspect where developments in the market raise a question over the wisdom of their inclusion in the EU AI Act are limitations on the use of emotional recognition technology in the workplace and in education.

Emotional recognition may become a standard feature of computing interfaces because it improves understanding and adds context. The release of ChatGPT4o on 13 May 2024, which includes emotion as a feature of its voice assistant, is illustrative. AI will need EQ as well as IQ to work well alongside us.

Risk of unintended consequences

The emotional recognition example has arguably not only failed to stand the test of time but might have had unintended consequences. Those who individuals who struggle with emotional recognition might have particularly benefited from such technology, and it surely was not intended to deny them such benefits?

Other examples of provisions that involve trade-offs and may have unintended consequences include the tiered approach to risk under the EU AI Act, and transparency and explainability provisions.

A tiered approach to risk sounds sensible and is preferable to applying overly restrictive regulation to low-risk applications. The problem is that its application involves the use of broad categorisations and is divorced from the detail of application specific regulation that, in combination with an application, determines risk.

High and low-risk activities can be expected to be included in a broad category designated high risk, and many high-risk activities may be low risk when account is taken of detailed regulation at the application level. An example is new medicines, which are high risk, but which must pass a highly regulated pathway to market. The expectation is that approved drugs are low risk, provided they are used as recommended.

The risk of unintended consequences arises because, by placing additional regulatory burdens on AI applications in broad high-risk categories, we prevent or slow their development. Returning to the medicines example, if AI tools that help identify candidate drugs face additional hurdles they may be utilised less, when clinical trials would have provided an assurance of safety in any case. Further such AI tools may also help predict efficacy and the risk of side-effects. Discouraging their development and use could harm safety overall.

Another example is autonomous transport where AI offers the prospect of risk reduction versus error prone humans. Such systems should have to demonstrate that they meet or exceed human levels of safety, but the development of the AI itself should not be hampered by additional requirements, since doing so could cost lives.

Turning to transparency and explainability, we are inclined to view these as unambiguously good, yet in practice they involve trade-offs: [46]

"Reflecting on historical and contemporary examples of scientific computing, we propose that the regulatory approach to AI should not be overly fixated on pre-emptive scrutability... We argue that strict adherence to explainable AI's

[46] Wareham, Romasanta, Priego and Osimo, January 2024, Regulating AI: lessons from scientific computing. *https://lisboncouncil.net/wp-content/uploads/2024/02/ESADE_Regulating_AI_Policy-Brief.pdf*

idealistic goals may indeed hinder the emergent capabilities of AI applications."

Trade secret protection is also relevant to enable and sustain innovation. Beyond utilising the public internet for training, identification of other specific data sets and methods of curation may in combination give away important trade secrets. Proprietary data sets (purchased, human curated and synthetic data) are also utilised, with synthetic data expected to grow in importance. Disclosure and reporting requirements could cut across the incentive to innovate to improve training data and models.

Details of model development, calibration and safety may also allow safety measures to be more readily circumvented or the models misused.

The EU AI Act is not a useful template for others, and they can afford to wait and see

The EU AI Act was developed in a specific institutional environment to address a specific challenge – harmonisation of regulation of AI across the EU single market – which may not be relevant elsewhere and which it in any case does not deliver. The application of AI is conditional on the rules governing applications, and the EU AI Act does not address these. Barriers to AI, and differences across member states, will remain under the Act.

Further, the EU AI Act includes general and specific provisions that overlap with existing regulation, which may not stand the test of time, and which may have consequences at variance with their purpose.

In relation to general purpose AI, the ideal approach would not be to regulate the models themselves. In relation to other AI, if the application involves risk, then one might expect it to be regulated by exiting regulation. If not, then the application, rather than AI, should be regulated.

The EU AI Act may, by adding an additional layer of regulation, will impede both the prospective productivity and safety benefits of AI, whilst doing nothing to address the challenge of aligning existing regulation and AI.

Some aspects of the EU AI Act may be harmful, both to achieving the prospective productivity and safety benefits of AI. Other aspects are at best premature and may never prove necessary as detailed law and regulation is adapted (a task that is in any case necessary). Finally, to the extent that new challenges do emerge over time (akin to rail safety following the development of stream locomotives) these are almost certainly not anticipated in the EU AI Act.

6. Way forward globally

Whilst AI may prove revolutionary, applications of AI will be regulated by existing law and regulation, and AI itself will be harnessed by business, consumers and regulators to address challenges.

We can therefore afford to wait and see before rushing to legislate for AI *per se*. Further, there is an opportunity to learn from implementation of the EU AI Act, which will depend on its detailed interpretation and enforcement over time.

Whilst the EU AI Act was conceived in a specific institutional setting, and whilst it is arguably already apparent that it is likely to have harmful unintended consequences, there is nevertheless a learning opportunity for everyone from observing the interpretation, implementation and impact of the EU AI Act over time.

The EU has an opportunity, given the scope for discretion and adaptation under the Act, to act pragmatically and minimise the risk of unintended consequences. However, the degrees of freedom are more limited than they are for others who have chosen not to rush to legislate.

Those with federal systems of government may need to prevent market fragmentation via inconsistent state rules. However, unlike the EU, they could do so whilst preserving a greater level of flexibility; indeed, federal law could seek to entrench flexibility over the use of AI, subject to existing context-specific regulation of applications.

Other countries have proposed or drafted AI related bills, including Brazil and Canada. They have an opportunity, before finalising legislation, to clearly define and limit the scope of AI regulation; to ensure that the risk of unintended consequences is minimised, including for safety where rapid AI adoption may contribute to safety overall; to limit overlap with other law; and to facilitate interoperability via mutual recognition with other jurisdictions.

Others, who may be contemplating their own law of AI, have an opportunity to pause and reflect on where effort would be most productively focussed; on premature legislation that risks holding back AI and may increase risk overall, or on the task of adapting existing law and regulation to align it with AI.

An immediate task in terms of adaption is ensuring that law and regulation permits AI wherever it is more efficient and/or safer than the *status quo* approach. Reform in the UK to permit autonomous ships is illustrative, however the agenda is much broader, in principle encompassing all areas subject to occupational and professional licensing.

Over time the need for other reforms to align regulation and AI, including potentially the revision of existing horizontal law, may become clear. The way in which disputes are resolved by the courts may also shed light on the potential need for legal reform.

We should also ensure that markets for inputs to AI, including electricity and data, are responsive to demand, consistent with limiting externalities and protecting privacy.

This approach is not novel, it is the one we followed with previous general-purpose technologies, adapting regulation over time in relation to issues such as train safety and data protection, but not rushing to enshrine everything in a law of steam or computing. This is the approach we should follow with AI.