The Politics of AI

Introduction

What are the politics of artificial intelligence? What does it mean when we talk about regulating the actions of a machine that expresses intelligence?

In planning our conference themes we have defined AI as a machine system capable of using data analysis to replicate human reasoning, both in achieving certain tasks and in learning from the results of, and adapting, its own protocols. This definition has allowed us to focus on four themes for politics and policy that run through each panel discussion.

The first is the question of what policy frameworks are needed for a world in which machine judgement displaces human judgement. When a machine advises an investor, dispenses a medical diagnosis, drives a vehicle or disseminates a news story, it is undertaking a regulated activity and navigating a human ethical landscape. What ethical challenges do these activities raise? Can we agree on a way to regulate them?

The second is the question of the impact AI will have on the workforce itself. The automation of human work is hardly new: AI is simply the next iteration in the long process of human labour displacement through technology. While unquestionably disruptive, such change is far from inherently negative. But its reach into a wide range of cognition-based jobs – a huge central strata of the modern services economy – will present policy challenges not experienced in the past.

The third relates to the data needed to make the deployment of AI in everyday functions possible. Learning how to think like a person can only be done by analysing the way people think. It is the aggregation for the first time of huge data sets of human behaviour and the computing power to exploit them that has unlocked the current potential in AI. But that data originates in real people, and how it is collected, protected and even remunerated – given its huge potential value – are important questions for policymakers to respond to.

A final theme is the industrial policy question that underlies this change. The race to lead the development of AI is shaping into an international one, but to what extent should it be the subject of industrial policy? To what extent are international frameworks for AI feasible or desirable? And how should policymakers balance support for a lead in the technology itself with adequate consideration of the policy challenges it is raising as it develops?

This conference is intended to debate these questions rather than answer them definitively. It is surely too soon to agree on answers to any of them. But, as this technology begins to transform the world of commerce and work, it is important to ask how it will also impact our politics, and demand sophisticated solutions from both businesses and policymakers.

Stephen Adams
Senior Director

What should a national strategy on AI policy look like?

The use of big data to drive machine learning and artificial intelligence will transform the way people live and work and the way economies function. In making a national strategy for AI, policymakers will have to weigh priorities that span everything from workplace regulation to industrial policy. What should a national strategy for AI look like? Is one even possible given the wide terrain and pace of change?

In the global race to dominate the creation of the technologies themselves, Europe appears to have made an uncertain start. China and US start-ups, for example, comprised 87% of equity funding to AI start-ups globally in 2017. The large markets of the US and China appear to have facilitated easy scaling. European responses through public funding have been isolated and focused on pre-commercialisation research in a way that has yet to register in international competition. The UK, which is allocating around €1 bn for AI development and benefits from a favourable AI ecosystem, is the most competitive in the race for AI funding, research and talents. Elsewhere in Europe, fragmented national efforts to attract inward investment or the €2.5 bn EU funding programme for AI seem unlikely to materially adapt this picture.

What conclusions should we draw from this? Is Europe's deficit reversible? Has the race to dominate these technologies sidelined debates about other aspects of this change in China and the US? And conversely - does Europe's relative weakness in this respect actually make it fertile ground for thinking through questions of competition, equity and regulation?

Under budgetary constraints, one potential strategy for European policymakers will be to turn to regulation to strengthen European global AI competitiveness, or to ringfence the EU market from external challengers. Liability policy is emerging as one playing field for the former: tax for the latter. Data policy will be pulled in both directions. Businesses may hope for a harmonised European landscape on liability to provide legal certainty and help investors weigh risk, but with European countries like Estonia developing their own liability frameworks, there is still a clear risk of fragmentation instead. On data, policymakers are on one hand asking whether EU data protection and privacy law is a check on ambitions for global competitiveness in AI. But on the other, they are tempted to see a more interventionist approach as a way to erode the dominance of US large tech firms, often linked to taxation issue.

Where should these balances between competitiveness and a sound and robust regulatory framework be struck? Can they be reconciled?

European policymakers have also been much quicker than their US and Chinese counterparts to adopt a defensive position on the potential disruptive social impacts of AI. Businesses have responded to this concern with a flurry of national AI codes of conduct and ethical guidelines. Competition policymakers have begun to think hard and publicly about conventional theories of market power and network effects. Yet with AI technologies still at an early stage of development, it is hard to predict their future implementations and long-term social consequences with accuracy.

Are attempts to get out in front of AI disruption the right approach? Are there parallels with globalisation in the lessons to be learnt about securing social acceptability for disruptive change? Do attempts to frame problems pre-emptively risk shutting down avenues for innovation or colouring public attitudes in a counterproductive way?

Work in progress: selected national AI strategies

	Industrial Policy	Data protection and localisation	Ethics/liability framework
United Kingdom 	£950m in public and private funding; AI funding via £1.7 bn Industrial Strategy Challenge Fund.	Some initiatives to encourage data sharing in public and private sector; open approach to data flows in trade agreements.	Establishing the Centre for Data Ethics and Innovation which will develop industry guidance and recommend regulation and legislation.
Germany 	Network of 12 R&D innovation centres; €3.5 bn in investment incentives over six years.	Creation of data sharing platforms; review of domestic competition framework; commitment to promote technologies respecting data privacy.	Federal government reviewing regulatory framework and "taking account" of recommendations from the Data Ethics Commission.
European Union 	€2.5 bn in Digital Europe Programme for AI, including testing and experimentation facilities; R&D funding available for AI research projects.	GDPR polices external flows; free flow of data regulation removing barriers within the EU; European Commission assessing link between anti-trust and data concentration.	European Commission's High-Level Expert Group on artificial intelligence developing ethical guidelines.
United States 	2017 defence budget commits $7.4 bn on AI, big data and cloud computing; 40% increase in government investment in AI R&D since 2015.	Open approach to cross-border flows; support for data flow provisions within trade agreements.	National Science and Technology Council prioritising "removing barriers to AI innovation"; ethical framework not included as a priority.
China 	Megaproject artificial intelligence 2.0; political support for Chinese tech giants' foreign investments in AI chips.	Heavy restrictions on cross border flows, e.g. the Internet Security Law; emergence of a debate over data protection.	Commitment to national data ethics framework by 2025.

AI and the transformation of European healthcare

Healthcare systems across the OECD, whether publicly or privately funded, face a range of challenges that threaten their financial and operational sustainability. The demographic crunch is the backdrop to this. Life expectancy has risen inexorably over recent years, but quality of life has often failed to keep pace. More people are living longer, but frequently with complex, chronic conditions that require ongoing treatments. Health systems have to meet this increase in demand with constrained resources, especially since the global financial crisis. Publicly funded systems in particular have received fewer resources than required as governments seek to economise, leading in some cases to a rationing of care.

Workforce shortages also threaten the sustainability of healthcare delivery. Although a more acute challenge in low and middle income countries, the inability of systems to recruit and retain sufficient clinical staff to meet rising demand could further harm patient outcomes. In the UK, for example, the health and social care sector has become increasingly reliant upon foreign labour, while employers often pay premium rates for temporary staff to fill gaps.

In its current configuration, the health system of a typical developed state risks becoming unsustainable without a substantial increase in resources, a clear reduction of expectations in what can be delivered for patients or a revolution in productivity. Although no silver bullet, AI could be part of the solution to these challenges.

The fundamental gains are in efficiency and cost-effectiveness. AI tools that can analyse medical data such as microscopic sections from biopsies can revolutionise the speed and accuracy of medical diagnosis. Advances in genomics and image analysis will also drive greater understanding of how complex diseases develop, allowing for more proactive and personalised treatment. Earlier diagnoses promise to reduce patient costs further down the line. AI solutions can make reporting more efficient, thus freeing up time for NHS staff to spend on direct patient care. Yet an underdeveloped policy framework and the presence of a risk averse culture amongst clinicians has slowed the deployment of these technologies – often rendering public healthcare providers less willing adopters of technology than they would like to be.

What are the primary obstacles to the deployment of AI in public healthcare? What simplifications of the regulatory and procurement framework would have the greatest effect? How are legacy technologies and commercial arrangements holding back change and how can this be addressed?

Many of these innovations require access to huge pools of patient data. Ensuring access to data for AI researchers and commercial organisations generally requires the breaking down of well-established institutional boundaries. Both requirements can encounter barriers in patient concerns over the sharing of personal data, as exemplified in the Royal Free/DeepMind controversy, and in the narrower commercial and institutional interests of healthcare providers, pharmaceutical manufacturers and tech companies. The greater use of AI driven diagnostic technologies also raises questions of trust and patient confidence. Whether those making algorithms should share the negligence liability shouldered by clinicians and whether regulators assess whether an algorithm is sufficiently robust in the same way as new drugs are evaluated are core questions.

What are the primary practical obstacles to data pooling and analysis inside public healthcare systems? How can they be resolved? Are there any unique aspects to data protection and patient trust in healthcare that need bespoke solutions?

Uptake inhibitors? AI in European healthcare

Key AI healthcare applications

- End of life care
- Research & training
- Well-being
- Early detection

Barriers to uptake?

- Navigating data protection
- Need for uniform formats
- Technology niches
- Lack of collaborative relationships, e.g. between insurers and doctors

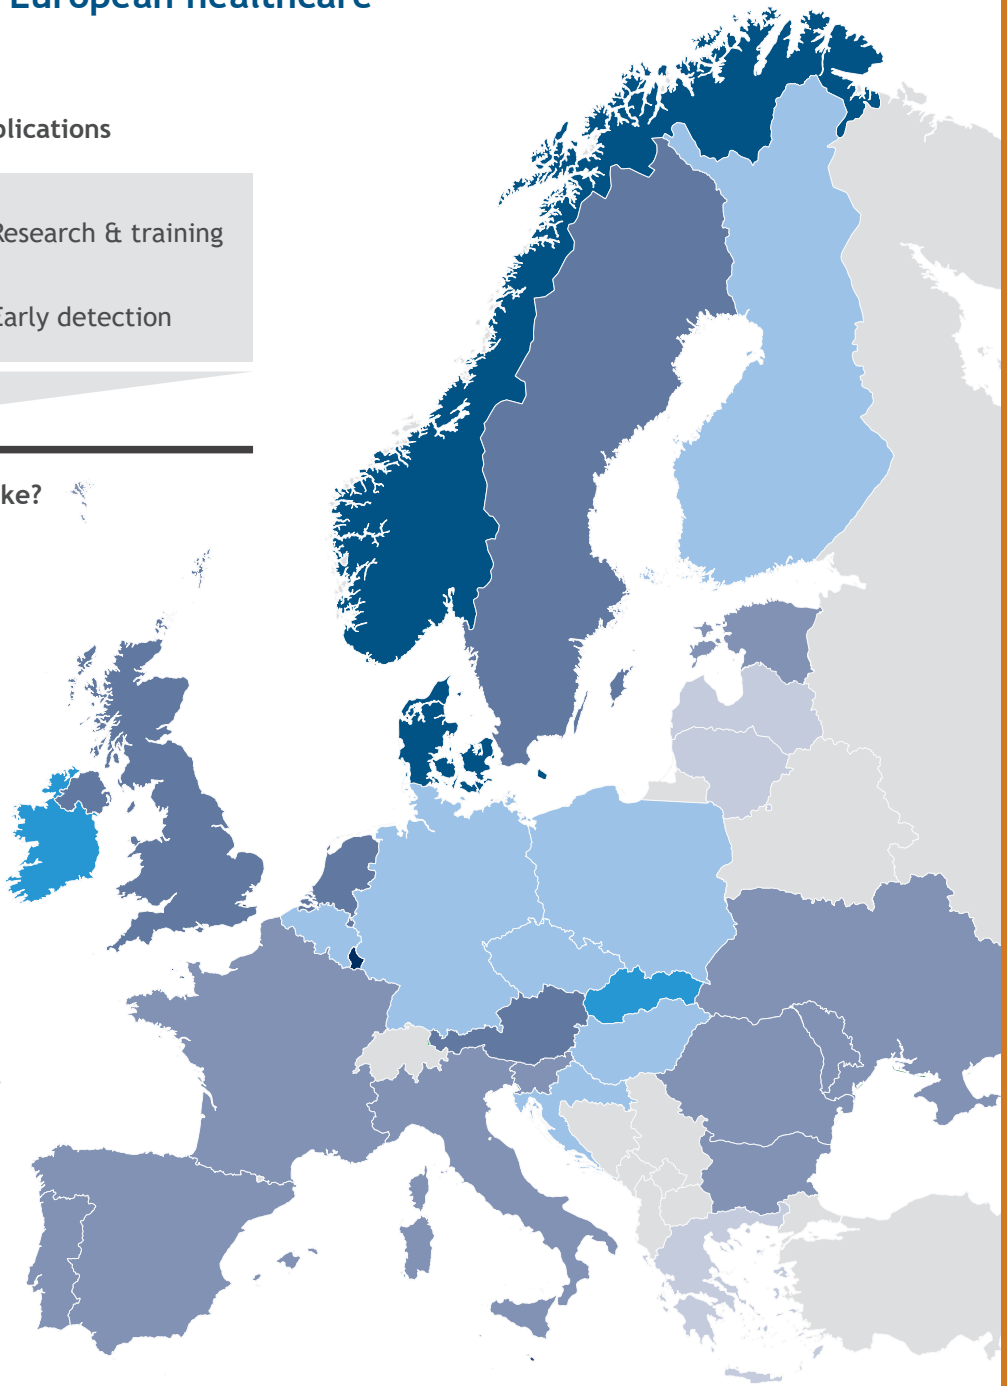

Risks in AI: re-evaluating the regulatory rulebook

AI is already transforming the delivery of financial services. This is true both at the highly sophisticated end of the market, where firms make use of AI to execute high-frequency trading strategies and at the level of the individual retail investor, where AI can tailor and deliver financial "robo-advice" to consumers. Deferring to machine judgement and automation is by no means new - automated trading is already the dominant form in most large markets. Yet AI clearly has much further to reach into financial decision making – from investment and trading strategies to client classification and product recommendations.

The scope in all of these areas to improve business processes, mitigate risk, lower costs and improve decision making is clear. From the predictive power of big data in insurance to data-driven 'contextual banking', customer information is increasingly utilised to drive improvements and customise services and refine risk profiles. This can create more personal, individualised services that consumers are familiar with – and increasingly coming to expect – in other areas of their lives such as shopping, entertainment or news provision. It can also lower costs like insurance premiums or investment fees.

But the same big data effects – machine learning and automated judgements – that drive these benefits also present potential challenges. Refined models of risk will inevitably mean rising costs for some – or even exclusion from services. Machine-learnt strategies may inherit the biases of the people who program and 'teach' them. Customers may also begin to resist the concept that the trade-off for lower prices is to "pay" in personal data, especially if this means a very close scrutiny of their financial, health or other personal data.

Will customers always be comfortable with the experience of being contextualised by their data? What questions does this use of data raise? And how will the inherent value of this data change the way it is stored, shared and protected?

The automation of consumer investment advice and personal financial management is also a big change for an industry that has always been based on human intermediation and which relies on trust in order to function. With lower barriers to entry, policymakers will hope that the increased availability of financial advice will boost savings and investment among an ageing population whose retirement provision is an increasing concern. At the same time, regulators will want clear lines of accountability and liability, to prevent mis-selling scandals from arising in the first place – and to know where to focus interventions if they do.

How should regulators and firms anticipate and address these risks? Is it possible to simply apply the existing rulebook to these technologies? What does a technologically neutral approach to regulating robo-advice actually mean?

AI in financial services

Firms

Sell

- ‘Segment of one’ – data analysis drives marketing and product customisation
- Enhanced pricing, clearer risk, lower premium, lower spread
- AI-enhanced KYC and AML e.g. insurers and doctors

Service

- ‘Contextual’ service – customer need anticipation and data-driven insights
- Robo-advice and manager disintermediation – automated investment strategy

Execute

- Automated execution
- Full automated trading strategies e.g. insurers and doctors

Shaping AI innovation

User concerns

- Terms of personal data ‘harvesting’, manipulation and protection
- Product exclusion
- Digital ‘intrusion’
- Distrust in disintermediation

Regulator concerns

- Liability for automated services
- Implications for market stability
- Consumer exclusion
- Consumer protection

Information quality in the age of AI

The initial promise of the internet age was for a free and open space to exchange information. However, the advent of algorithm-driven content raises important questions about what we see, how much we share a common information landscape and the reliability of what we rely on for news.

One of the biggest transformations of media is the data-driven turn towards users and audiences and the advertising model that governs it. This model relies heavily on speed and volume of delivery: the quality of information is less important than whether it is read. Platforms also compete by customising news, entertainment and information based on preferences 'revealed' by clicks, likes and searches. Users' data is in turn sold to advertisers, whose interest is generally in user eyeballs rather than the quality, breadth and objectivity of the information they are seeing.

Positively, digital media has been shown to expand the number of sources to which a reader is exposed – contradicting oft-expressed concerns over 'echo chambers' and 'filter bubbles'. Citizens' media literacy is developing to account for a wider range of views. However, the current social media landscape poses other problems: the competition for attention encourages inaccuracy and extremism. This continues to shape the public sphere into a place where emotion and simplicity are the norm – hampering efforts to persuade through nuanced argument. This has important implications for politics and political campaigning.

The rise of AI-driven disinformation is also threatening to undermine our sense of shared reality. Synthesised text, audio and video powered by machine learning threaten to produce artefacts indistinguishable from real life, both stimulating confusion and providing plausible deniability to gaffe-prone or scandal-hit politicians. Combined with the scaled sharing platforms we use, including encrypted messaging, this is already creating enormous social discord in countries like India and Myanmar.

But AI can also be used for enormous social benefit in the information environment. Data aggregation and analytics tools similar to those used by intelligence services and law enforcement could create efficiencies and opportunities in journalism itself, as well as business processes such as regulatory compliance, know-your-customer checks and anti-money laundering efforts.

Can an advertising revenue model for large platforms and information aggregators be made more conducive to heterogenous sources of information and news? If so, how?

Transparency is a big issue. Social media algorithms are constantly evolving and increasing in complexity, turning into mysterious black boxes of decision making. Users have little understanding of what personal data is used and how and even less control or ways to influence these processes. What platforms tell their advertisers they know about their users can be somewhat at odds with the picture they present to users themselves.

How much transparency is required with respect to the use of user data on platforms? Where should the line between proprietary analytical technology and user interests in their own data and its collection and interpretation be drawn? Does the advertising driven model create incentives to use data in ways that could be different in other revenue models?

Revenue and AI information curation models in the platform economy

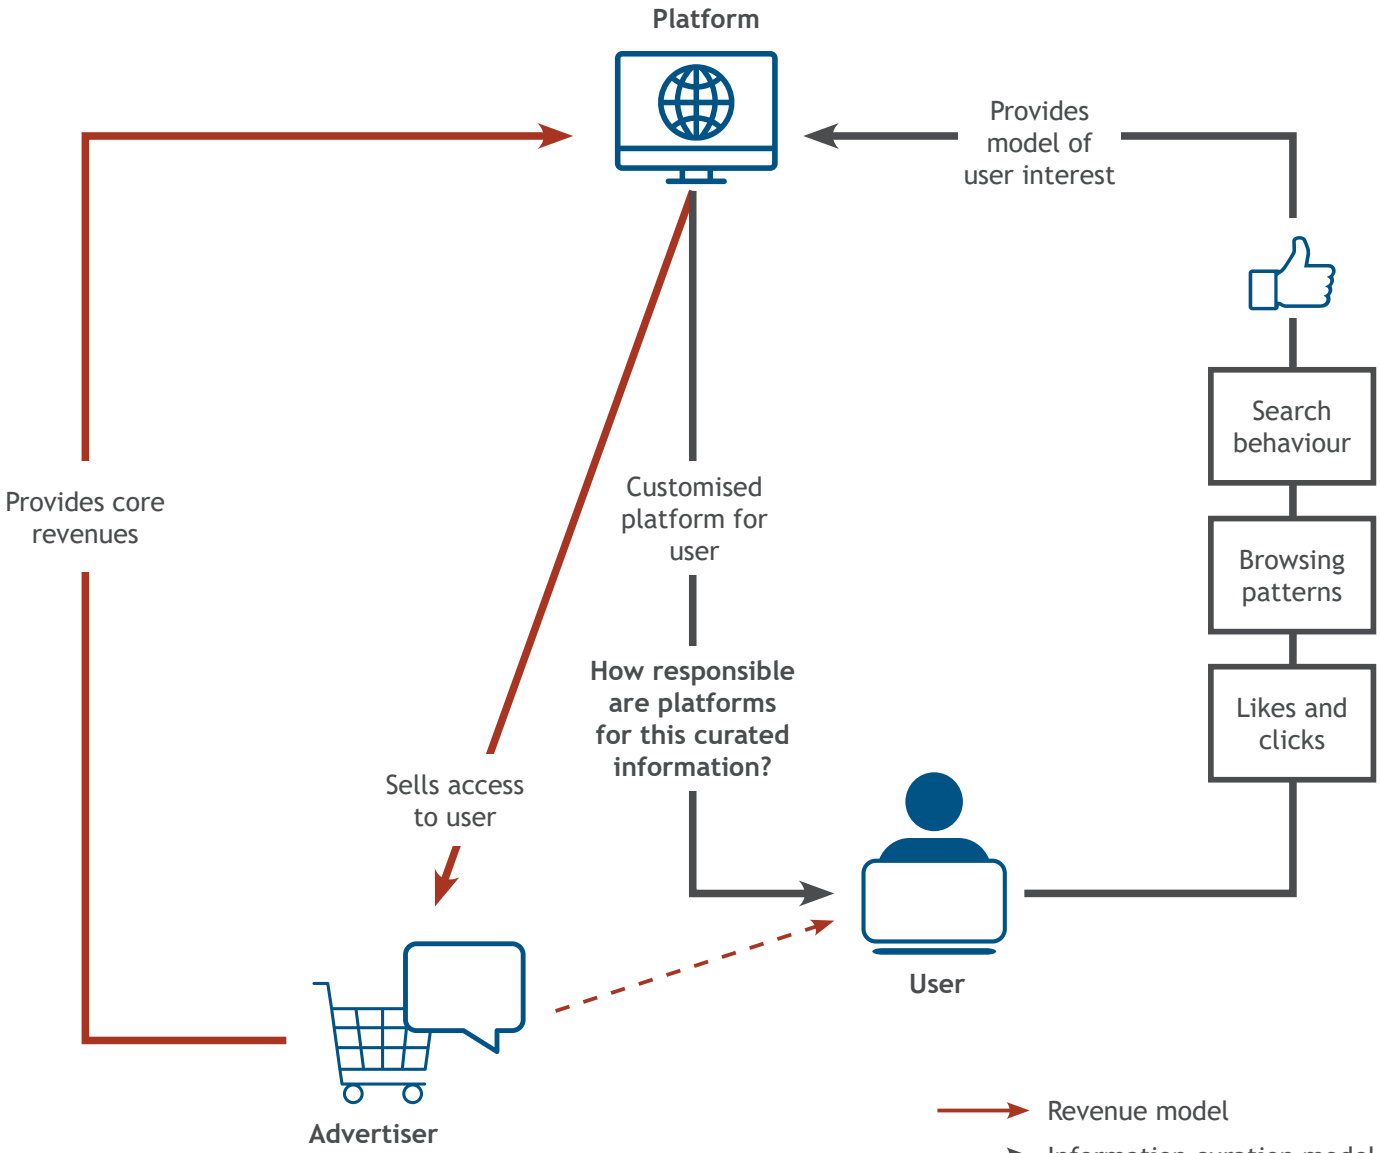

The economics of data

Most assessments of the impact of AI on European labour markets conclude that it will have two key impacts: displacing human labour and driving productivity. In the long run, both of these effects can be positives. In the short run, they are likely to be disruptive and disorienting. The reach of automation into parts of the cognitive labour market that have heretofore been protected from it may be a social shock, and raises the prospect of a hollowed out labour market where elite cognitive labour and non-routine low-cost manual work sit above and below a disappearing landscape of traditional skilled middle class jobs in both manufacturing and services. While new roles are likely to be created as well as old ones displaced by machine judgement, the detail of this evolution is inherently uncertain in advance.

Concerned policymakers are inevitably beginning to worry about the disrupted. Their worries tend to focus in some key areas. One is the skills demands that will come with this transition. Facilitating the development of new skills and enabling workers to retrain are a constant focus, perhaps especially in manufacturing. But the role of government and businesses in delivering on this remains undefined.

How much of this burden can and should be borne by the private sector? What is the public sector's role in a model in which the most intensive retraining needs are likely to be individuals over 50 and long out of the public education system? Beyond STEM skills, what other kinds of skills needs is the AI revolution likely to create in areas such as communication?

How to make the labour market sufficiently adaptive to restructure around these changes without sacrificing worker security is another inevitable challenge. Social welfare systems would struggle with large scale job displacement – but employee protection could be a check on effective deployment of technology.

If there is a balance to strike here, where is it? How do we ensure that worker protection does not in itself become an incentive to automation in a self-defeating way? What challenges of labour mobility need addressing?

Some observers have gone much further in their thinking and asked if the new economy is potentially redefining the nature of labour itself. Labour productivity improvements are expected to account for over 55% of all GDP gains from AI between now and 2030. This means a further fall in labour's share of GDP. Some argue that in an economy driven by data, the personal data that individuals provide to drive the operation and profitability of the era's defining technologies is analogous to their labour – and they should be remunerated for it, especially if these technologies are simultaneously devaluing their physical or cognitive labour in other ways.

How credible are these arguments? How feasible is the idea that people can and should be rewarded for the work their data does and the profit it earns for others?

Cognition as work: the reach of AI into current labour markets

OECD 2016: 14% of jobs highly automatable + 32% at significant risk of automation

McKinsey Global Institute 2017: 30% of work activities could be automated by 2030

Frey and Osborne (Oxford) 2017: 47% of total US employment at a high probability of automation by 2033

Bruegel 2014: 54% of EU jobs at risk of 'computerisation'

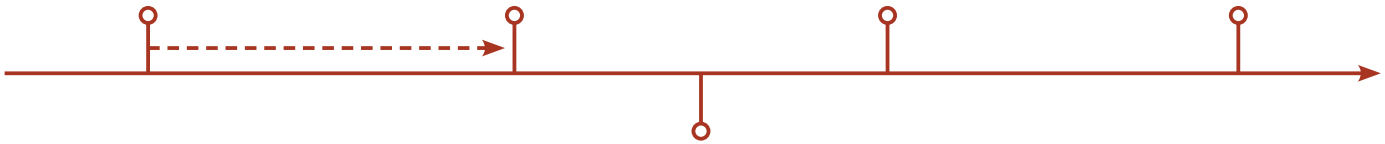

How far will AI reach into information processing, communication and empathy tasks?

Process	Synthesise and analyse	Communicate	Convince	Advise	Care for

What will define the politics and policy of this transition?

Managing political backlash?	Demand for adaptation skills?	Demand for social welfare redesign?	Rising inequality and its implications?

About Global Counsel

Global Counsel advises international investors and businesses on political and regulatory risks in markets around the world. We help investors and companies in a wide range of sectors to anticipate how government policies and regulations will impact on their investment plan or business strategy and to develop and implement responses to these challenges. Global Counsel has offices in London, Brussels and Singapore.

www.global-counsel.co.uk

London

5 Welbeck Street, London W1G 9YQ
E: info@global-counsel.co.uk
T: +44 (0) 203 667 6500

W: www.global-counsel.co.uk
@global_counsel

Brussels

Résidence Palace, Rue de la Loi 155
Brussels 1040
E: info.eu@global-counsel.eu
T: +32 (0) 2808 5250

Singapore

1 Wallich Street, Level 14-01, Guoco Tower
Singapore 078881
E: info.asia@global-counsel.com.sg
T: +65 6403 3877