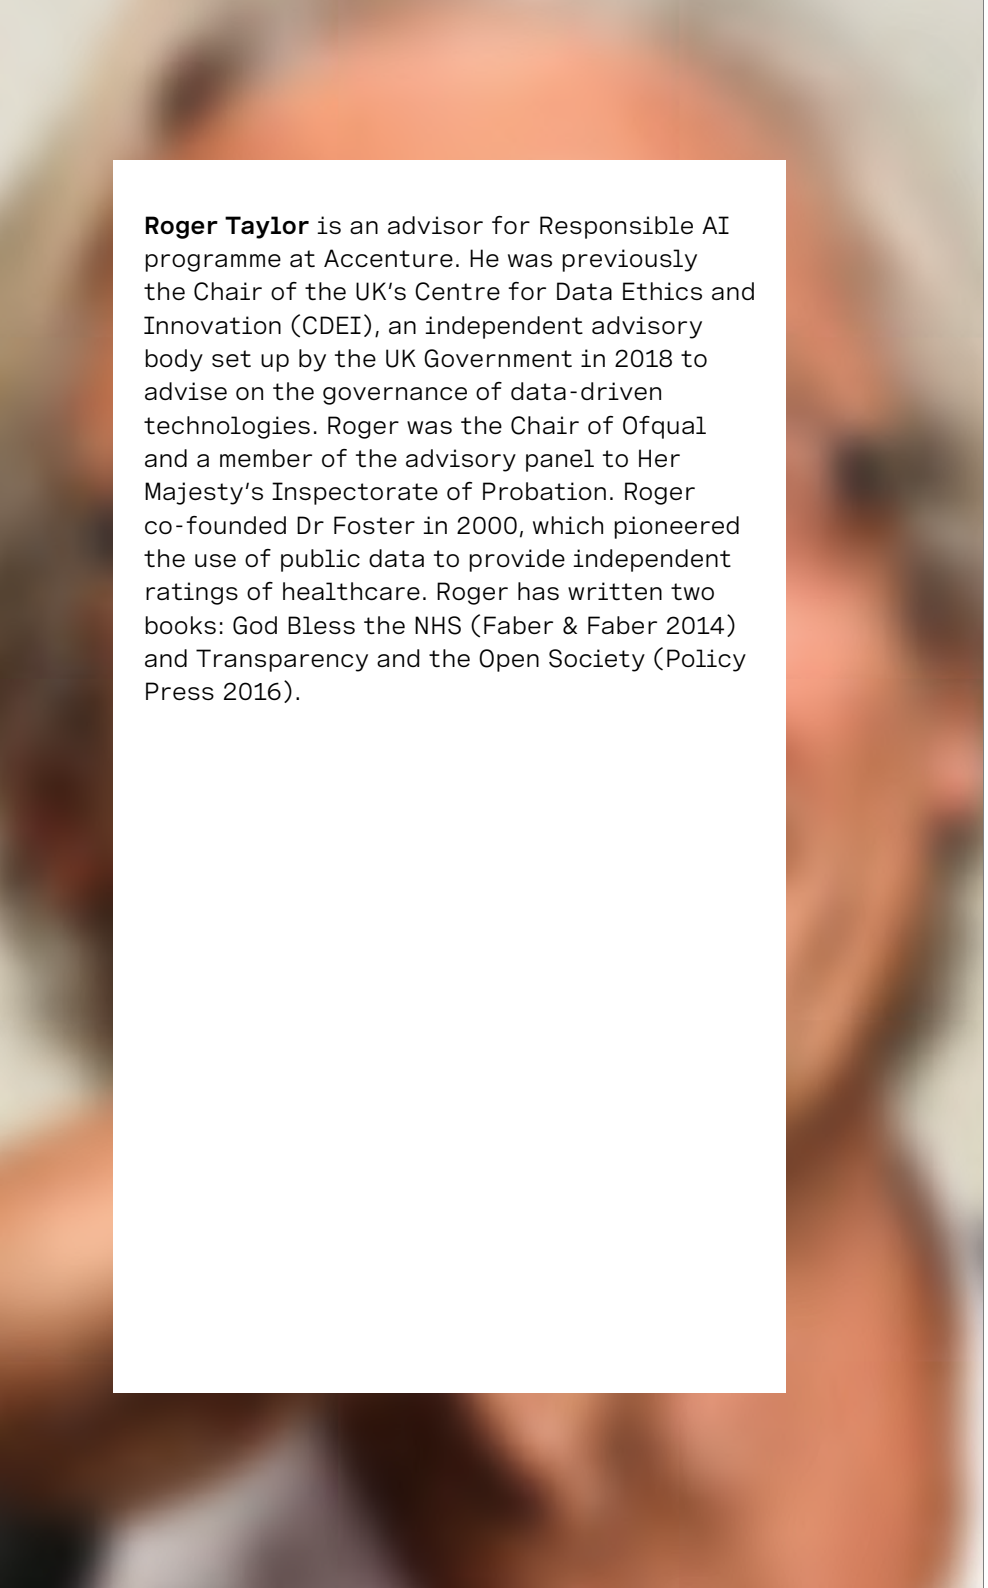

Roger Taylor is an advisor for Responsible AI programme at Accenture. He was previously the Chair of the UK's Centre for Data Ethics and Innovation (CDEI), an independent advisory body set up by the UK Government in 2018 to advise on the governance of data-driven technologies. Roger was the Chair of Ofqual and a member of the advisory panel to Her Majesty's Inspectorate of Probation. Roger co-founded Dr Foster in 2000, which pioneered the use of public data to provide independent ratings of healthcare. Roger has written two books: God Bless the NHS (Faber & Faber 2014) and Transparency and the Open Society (Policy Press 2016).

New approaches to data stewardship in education

Roger Taylor, Open Data Partners

At the heart of problems created by the digital economy lies control of data. Harms arise when an organisation has exclusive control over data about people, but has interests that diverge from theirs. Proprietary control over data, of the sort exercised by large platforms, also hinders innovation. It allows companies to achieve market dominance and to crush competition. Changing the way that data is controlled is a way to enable innovation while also providing a level of protection that current rights-based data regulation is unable to secure.

The current data protection regime is unable to provide adequate protection. Individual consent does not provide a way to understand how data about me is used. To truly understand this, I would need to know the impact of that data used, compared to other people. The regulatory focus on purpose - an important and valuable protection - is of limited use in today's data-driven world where people want personalised digital services but are at risk of being manipulated or discriminated against. My problem is not stopping people using my data for purposes I disapprove of; my problem is stopping people using my data for purposes I approve of, but doing it in a way that is ultimately damaging.

Look at education, for example. The use of data to drive

decisions and tools in education is potentially of enormous benefit, but it could also cause serious harm. The problem is not deciding whether to say 'yes' or 'no' to a particular purpose; it is knowing how we can protect ourselves from the poor use of data for purposes that are theoretically beneficial but prove, in practice, to be harmful.

This is why many technologists advocate a much more radical solution - breaking up the inherently monopolistic characteristics of digital markets by separating control over data from the provision of data-driven services. This would allow for unconflicted data stewardship organisations to monitor how data is being used, represent the interests of data subjects, prevent misuse of data and ensure appropriate levels of competition.

This is the principle behind the idea of data trusts, which has attracted many champions from technology industries because it offers a uniquely compelling vision of how to address problems in digital markets.

One barrier to the adoption of these approaches is the belief that data protection regulation can solve a problem it is not designed to solve (see Figure 1). We should be sceptical of regulatory solutions that rely on data protection regulation and do not take full account of how this approach could be applied in practice.

An example: Google Classroom

Google provides equipment and digital services free to schools, such as email or apps for setting and receiving assignments. Google says the data collected is used only to monitor and improve these services. Children can, with parental consent for under-13s, also use the Chrome browser or Google maps. If they use Chrome they will see adverts, but Google says no data from children is used to personalise these adverts.

Google's suite of education products has been criticised from many angles (e.g., Krutka et al., 2021). In a feature in Fast Company last year, it was accused of disguising its business model, 'making it almost impossible to ascertain what data it collected about students and what Google uses it for' (Williamson, 2021).

Last year New Mexico started legal proceedings against Google, claiming it was illegally tracking the online behaviour of children under 13. The case was initially dismissed, but New Mexico appealed and Google settled. They admitted no breach of the law but agreed to do more to police age-screening on the app store and to fund an education initiative in the state.

The claim from New Mexico was that Google had collected information without getting clear consent and concealing its activity. The complaint accused Google of 'infiltrating' schools; of claiming its product was free when in fact it 'comes at a very real cost which Google purposefully disguises'. It said the company was 'mining children's data' for commercial benefit.

After the settlement New Mexico Attorney General Hector Balderas said: 'There are incredible risks lurking online and we should do everything we can to protect the privacy of children'.[1] This is true. However, it is not clear that his actions have had any significant impact on these risks. Google may not be breaking the law, but it has not become any easier 'to ascertain what data it collected about students and what Google uses it for'. To the extent, if at all, that Google has been infiltrating schools and imposing a 'very real cost' on children, nothing of significance has changed. One commentator called it 'fundamentally a victory' for Google, and pointed out that a new Google-branded education institute in New Mexico was a win for the business (Gold, 2021).

While Google is no doubt gathering a large amount of data, which it will use for commercial benefit, and which may harm people, the problem is that data protection is an ineffective tool to combat this risk. Attacking Google for unauthorised data use or inadequate consent misses the target. Even with all the necessary consent and legal authorities in place, the risk that the data is used in a way that harms young people remains. The question is not whether the company has the legal authority to use the data; it is whether it is doing so in a way that is harmful or beneficial.

For example, companies such as Google will typically establish a legal basis that allows them to use data to improve their service. It makes little sense to object to this in principle, but 'improving services' could mean something relatively innocuous, such as designing better ways to present email. It might also mean using the data to build artificial intelligence

(AI) that reads children's essays, monitors the speed with which they write, sees what time in the evening they do their homework and starts to build an understanding that could inform recommendations to teachers based on highly personal profiling. The second of these might be an enormously beneficial thing to do, but equally, it might be extremely damaging. As things stand, we have little way of knowing whether Google's work to improve its products is innocuous, brilliant or destructive.

The need for a new standard of practice in data driven systems

This type of market failure is not new. There are many products where the market would not work without quite specific regulations regarding information. Medicines are one. You cannot tell whether a pill works by looking at it. And it is not safe to find out by trying it. Instead, we have an elaborate regulatory mechanism that sets standards for how information is generated and shared to assess a product's efficacy. Cars and airlines are similar, in that you are safe to choose a car on the basis of its shape or an airline on the basis of its food because regulation does the work of ensuring the wrong choice is unlikely to cause serious harm.

AI and other complex data-driven systems are similar in that the quality of the product or service can only be assessed with knowledge and quite specific datasets.

Data-driven systems present two additional challenges. First, it is difficult to tell in advance where the dangers might be. No one imagined that using machine learning to build recommendation systems in social media would help unleash a pandemic of misinformation. Second, we are not talking about one class of products - it is a fundamental technology that is altering a wide range of products and services, introducing new risks to all of them.

Data protection law was developed to control the purposes for which data is used, and is grounded in the concerns that arose during the initial development of databases. We face very different risks today that cannot be effectively addressed with data protection law.

What would a new approach to data governance look like?

A number of different elements have a role in reshaping digital markets, but the key ideas are:

- Separation of the data layer from the application layer in the architecture of digital services
- Independent governance and control of the data layer by organisations that are legally excluded from providing apps and services and that have duties towards data subjects (e.g., data trusts)
- Protection of individual data rights through personal data stores (or rights to data portability and reporting).

These ideas are independent and there are examples of each. For example, Open Banking in the UK is a mechanism to enforce portability of financial data;[2] the medical research field has a number of 'data governance' bodies that oversee access to data (e.g., HDR UK); and in the commercial area, shared data pools such as 'Skywise' allow companies involved in building Airbus aeroplanes to manage access to a shared pool of data.[3]

However, the biggest opportunities lie in bringing these ideas together and applying them to the provision of personal digital services. Together they create a virtuous circle that can support a market for demonstrably beneficial innovation. They allow for decentralised management of digital IDs and create the space for a market of 'digital agents' who represent the interests of individuals and communities, enabling individuals to maintain control over how data about them is used while at the same time empowering organisations capable of turning these rights into effective market or regulatory power.

Separating the data from the application creates a market incentive to drive the adoption of data standards to the extent that services and apps make use of common underlying data. This allows for greater competition in the provision of these services. It also allows for external experts or regulators to develop the skills to interrogate and interpret the data that are equal to those of the organisations providing services.

Establishing separate governance for the data layer means

that data users have to make the case for their use of data to an organisation that is its equal, both in terms of its ability to control data and to understand how data is being used.

The detail of how these mechanisms are best applied in any particular area depends on the context. However, if these arrangements were applied in education, we can imagine a scenario in which a personal data store would hold a defined set of information about the pupil, including the data generated by the school. The school would then operate a data trust on standard terms with pupils or parents and caregivers. Such trusts might be federated across similar schools and operated by an independent trustee body. The terms would set down not just the purposes to which data can be put, but also the way in which the impact of data use is assessed and the mechanisms by which data subjects are kept informed and able to exercise choice.

The trust would also set out the terms on which providers of digital education services could access data. For example, the trust might set requirements for data must be returned to the data layer (e.g., activities, test scores) in either standard or proprietary formats. Such arrangements might require, for example, that any assessment of bias or benefit would be based on data held in the independent data layer, not on the provider's own data systems.

In effect, this mechanism replaces regulation with market incentives. This can then ensure the appropriate level of resource going into these activities - activities that are value-creating for society and the economy, but that would likely be prohibitive if framed as a regulatory requirement.

This approach would end the pretence that individual consent is enabling people to exercise meaningful control over data use. Instead, agents acting on behalf of parents and caregivers, children and schools would have the powers and capabilities necessary to protect their interests.

Such an arrangement would also afford greater freedom to providers of digital services to innovate and improve services, without increasing risks of data misuse.

Why is there limited progress towards reforming digital markets?

Progress towards creating this new world is not due to lack of enthusiasm or hard work. For many years leaders in the technology industry have been calling for root-and-branch reform of the data economy and working to achieve it.

In 2021 Tim Berners-Lee launched Inrupt, a company that builds on the work of the SOLID data standard for personal data stores, recognising that we need to rebuild the data economy from the ground up. In the UK, Professor Irene Ng has a similar initiative, HatDex, which enables individuals to require their data to be held in a separate database that they own.

The Open Data Institute has championed the use of data trusts to create a new layer of governance over data use.[4] Neil Lawrence, former Director of Machine Learning at Amazon and now DeepMind Professor of Machine Learning at the University of Cambridge, has established the Data Trust Initiative to support the implementation of such arrangements (Gardner, 2020; see also Delacroix & Lawrence, 2019). The Mozilla foundation has also been active in encouraging new approaches to data management.

Despite this, these ideas receive insufficient attention in discussions about regulating digital services, whether in education or in any other area.

Defend Digital Me, which campaigns on the use of data in education, has made detailed recommendations to prevent abuse of data, but does not address the need for wholesale reform of the relationship between control over data and provision of data-driven services (Defend Digital Me, 2020).

The UK Government, which advocates a strongly pro-innovation stance towards data (DCMS, 2021), has been very clear in setting out how it intends to reform data protection to remove regulatory barriers to innovation. In comparison, its comments on data stewardship lack detail and substance. The consultation proposals for reform of the General Data Protection Regulation (GDPR) had extensive analysis of problems with current regulatory arrangements, but little to say about new approaches to data governance.

There are several things that can account for this. The first is that the problem is hard. It is hard for policymakers to get

their heads around the many difficult questions any implementation of these new arrangements raises. For example, to what extent is it necessary to impose a minimum level of custodianship on a market, or should this be something that individuals or market participants can opt in to? The second option is more appealing to governments because it requires less action, but is also less likely to succeed.

If a minimum standard is imposed, what are the legal mechanisms that are best suited to doing this? What institutions are required to oversee this? How far would its powers extend in setting standards and/or requiring data sharing with end service providers or intermediate data agency /data trust services?

A second problem is that the answers to these questions are very context-dependent. They would vary, for example, according to whether the service under discussion is safety-critical, highly regulated, state-provided, foundational (e.g., identity) or an entirely optional consumer service.

It is easier for governments to set out frameworks and overarching mechanisms. The EU is implementing exactly this sort of approach through its Data Governance Act that establishes the basis on which data-sharing mechanisms might operate.† The UK has similarly been exploring 'enabling' frameworks to allow for such mechanisms to exist.

However, the market failure that makes new forms of data stewardship necessary is the same market failure, which means that simply 'enabling' solutions to exist will be insufficient. The role of government here is not to enable, but to deliver.

The last problem for governments is uncertainty. The complexity of the situation, and the range of possible solutions, means that there is no way to reliably and comprehensively design such a complex set of arrangements in advance. This is one of a class of problems in which the solution can only be identified by first making a commitment to put new arrangements in place, and then working through the issues with stakeholders. Such situations are not unusual in life, but they are never comfortable for governments for whom the appearance of control is so vital. The necessary engagement from interested parties will not be available without a commitment to implementation and funds to support the work. The only way to make progress is to recognise that reform is essential, decide where to implement reforms and have the political courage to commit to seeing it through. 'Commitment' here would mean establishing a competent legal authority to oversee reforms, giving it an appropriate budget, and setting a principles-based framework within which to operate along with target dates and reporting requirements.

This can be challenging and off-putting to government that may lack the skills and knowledge to feel comfortable about the risks. Elected representatives currently face no compelling reason to wade into such difficult waters. These policy ideas do not offer quick solutions. They require long-term strategic planning and significant investment to build digital services for the next generation that are trustworthy and beneficial. Politicians are happier leaving it as a nice idea and offering warm words of encouragement.

Applying new data stewardship models in education

If a government were to demonstrate the necessary vision and courage, education is an interesting and promising area where intervention to reshape the digital economy could bring significant benefits. The market is not one that has been staked and claimed - in the sense that there are no dominant education-specific services that generate their value from the proprietary exploitation of the knowledge contained within mass data collection. There is significant potential benefit from the use of AI and data-driven technologies if done right, and there are significant risks in leaving it to current market arrangements. There is widespread acceptance that government has a key role in assuring the quality of education, including digital education services.

Crucially in education there is a credible route to success. Moving to a new model is much easier if there is an 'on ramp' of deliverable benefits that start at a low level and build as a system develops. In education, relatively simple steps - such as giving people digital certificates for their qualifications -

† This is not a criticism of the EU as it would be difficult for an overarching body such as the EU to do more than this. The criticism is of national governments that could and should do more.

provide both a useful service (you do not have to find paper certificates for job applications) and creates the basis for the establishment of new data stewardship arrangements based on personal data stores.

The tools we need to move beyond the current debate about data protection and instead initiate a discussion about reform of the data economy are available to us. However, it needs the catalyst of political pressure and political will to change the way in which this market operates. Without it, we will not be able to deliver a safe, innovative, digitally supported education system.

Figure 1: GDPR, data protection and harmful data-driven services

A key assertion in this essay is that data protection laws as currently constructed cannot offer adequate consumer protection. This claim could be supported by examining the extent of its impact on the behaviour of the large data platforms. However, given limited space, it is simpler to ground in some more fundamental observations about the mechanisms of data protection (for a much longer discussion of these issues, see Taylor & Kelsey, 2016).

The core principle of data protection law (and a good principle, too) is that data should not be processed without lawful grounds. This gives consumers and regulators the power to 'pull the plug' and halt data processing. Consumers can do this by withholding consent, regulators by rejecting the legal basis of processing.

Although data privacy advocates recommend 'pulling the plug', this recommendation cannot protect consumers for the following reasons:

1. *People want personalised services*. If people were willing to do without, then the power to pull the plug would fix the problem. However, the majority of people in the UK, USA and Germany are in favour of the use of personalisation in a wide range of applications including recommendation systems and advertising (Kozyreva et al., 2021). People in the UK are also in favour of using personalisation to make recommendations in education by identifying educational needs (CDEI, 2020).

2. *Regulators do not have the powers and capabilities necessary to identify and address harmful personalisation*. People need protection, not from personalisation, but from personalisation that is biased, manipulative or harmful in other ways. You cannot achieve this by saying 'no' - partly for the obvious reason that rejecting things does not force people to give you what you want - but more importantly, because you first need to be able to

identify whether or not a particular use of data is harmful. This is rarely immediately apparent. For example, to identify bias you need to look at how a system treated large groups of people and compare it to similar systems.

3. *The scale of the task makes it implausible that a single regulator could not have the powers and capabilities necessary to identify and address harmful personalisation*. Europe has proposed a new AI law to address the problem of harmful decision-making. It does this by creating an obligation on users of AI systems to demonstrate they have a system in place to manage risks. This recognises the purely logistical problem of having a single regulator attempting to oversee the fairness of data-driven decisions in health, education, finance and life in general. That is an unfeasibly large task, and the issues concerned are, in many cases, covered by other existing legal obligations.

4. *The best way to protect the consumer is to ensure that someone other than the provider of a data-driven service has the capability, the power and responsibility to assess whether it is beneficial or harmful*. To assess the fairness, accuracy and/or harmfulness of data-driven systems requires an ability to compare between systems and to look beyond the immediate data on which the system runs (which, in the main, will confirm the vendor's view of the system). The harm that can come from misuse of data in education is that it hinders education or negatively effects children in ways that are not obvious to those using such systems. An assessment of whether something is harmful will depend crucially on the ability to compare it to alternative approaches to education and understand the impact in the wider context. The creation of data agencies, data trusts or regulators that manage shared data pools within key industries provides just such a mechanism.

CDEI (Centre for Data Ethics and Innovation). (2020). CDEI review of online targeting

Day, E. (2021). Governance of data for children's learning in UK state schools. Digital Futures Commission, 5Rights Foundation

DCMS (Department for Digital, Culture, Media & Sport). (2021). Data: A new direction. 10 September

defenddigitalme, (2020). The state of data 2020: Mapping a child's digital footprint in the state education landscape in England

Delacroix, S., & Lawrence, N. D. (2019). Bottom-up data trusts: Disturbing the 'one size fits all' approach to data governance. International Data Privacy Law, 9(4), 236-252

Gardner, R. (2020). New Data Trusts Initiative will spearhead community-focused data governance. Department of Computer Science and Technology, University of Cambridge, 21 October

Gold, A. (2021). Google settles children's privacy suits brought by New Mexico. Axios, 13 December

Kozyreva, A., Lorenz-Spreen, P., Hertwig, R., Lewandowsky, S., & Herzog, S. M. (2021). Public attitudes towards algorithmic personalization and use of personal data online: Evidence from Germany, Great Britain, and the United States. Humanities and Social Sciences Communications, 8, 117

Krutka, D. G., Smits, R. M., & Willhelm, T. A. (2021). Don't be evil: Should we use Google in schools? TechTrends, 65, 421-431

Livingstone, S., Atabey, A., & Pothong, K. (2021). Addressing the problems and realising the benefits of processing children's education data: Report on an expert roundtable. Digital Futures Commission, 5Rights Foundation.

Taylor, R., & Kelsey, T. (2016). Transparency and the open society. Policy Press.

Williamson, B. (2021) Google's plans to bring AI to education make its dominance in classrooms more alarming. Fast Company, 28 June.

1 New Mexico press release: https://www.nmag.gov/attorney-general-hector-balderas-announces-landmark-settlements-with-google-over-childrens-online-privacy

2 https://openbanking.org.uk

3 https://aircraft.airbus.com/en/services/enhance/skywise

4 https://theodi.org/project/data-trusts