

Data management and use: governance in the 21st century

Priorities for data governance: discussions at a British Academy and Royal Society seminar on 16 October 2017

THE
ROYAL
SOCIETY



BRITISH
ACADEMY

for the humanities and social sciences

The text of this work is licensed under the terms of the Creative Commons Attribution License which permits unrestricted use, provided the original author and source are credited.

The license is available at:
creativecommons.org/licenses/by/4.0

Images are not covered by this license.

Contents

A: Introduction	4
B: Priorities for data governance	5
Discussions at a British Academy and Royal Society seminar on 16 October 2017	
C: Seminar papers	10
Michelle Brook and Anthony Zacharzewski	
The Democratic Society	11
Simon Burall	
Involve	13
Estelle Clark	
The Chartered Quality Institute	15
Guy Cohen	
Privitar	17
Matt Fenech	
Future Advocacy	19
Tim Hill	
Law Society	21
Jackie Hunter	
BenevolentBio	22
Irene Ng	
HAT Foundation Group	23
Nora Ni Loideain	
Information Law and Policy Centre at the Institute for Advanced Legal Studies	27
Kieron O’Hara	
University of Southampton	28
Hetan Shah	
The Royal Statistical Society	30
Amanda Smith, Jeni Tennison, Leigh Dodds and Peter Wells	
Open Data Institute	31
Maeve Walsh	
Corsham Institute	34
Marko Balabanovic, Michele Nati, Anat Elhalal, Phil Young and Paul Galwas	
Digital Catapult	36

A: Introduction

As data collection activities continue to increase in speed, scale and variety, and the analytic techniques used to process these datasets become more sophisticated, individuals and communities are affected in new and unexpected ways. The Royal Society and British Academy's report *Data management and use: Governance in the 21st Century* addressed this changing data landscape and recommended a principled approach to data governance, and called for stewardship of the entire data governance landscape.

The Academies hosted a seminar to explore the priorities across sectors for such a stewardship body. The discussions at the seminar are summarised in this paper, and set out a set of governance needs, practical challenges and conceptual concerns that any such body could take on. Since the seminar, the Government has announced in its budget statement on 22 November 2017 the creation of a new Centre for Data Ethics and Innovation to enable and ensure safe, ethical and ground-breaking innovation in AI and data-driven technologies. The Nuffield Foundation has been working closely with a number of organisations, including the Alan Turing Institute, Royal Society, British Academy, Royal Statistical Society, Omidyar Network, TechUK and the Wellcome Trust, on plans to establish an independent Convention to tackle the ethical and social issues arising from data use, artificial intelligence and associated technologies.

The summary below is followed by a set of papers submitted ahead of the seminar which set priorities for data governance. We hope these considerations will inform the important deliberations of these bodies.

Disclaimer: This is a note summarising the discussion and debate at the British Academy and Royal Society's workshop on *Data management and use: governance in the 21st century*. It is not intended to represent the views of either the British Academy or Royal Society, nor does it represent the views of individual attendees at the event.

B: Priorities for data governance

Discussions at a British Academy and Royal Society seminar on 16 October 2017

Changing data, changing society – the challenges of data governance

Things are moving fast. Data is everywhere and increasing in an uncontrolled way. This influx of data creates exciting opportunities for improving pretty much every sector in public life and in private industry, but at the same time it generates concerns. There are huge new opportunities but also huge new challenges.

We need to make sure that we are entering this new world in a way that includes everybody; so that the benefits and the risks are equitably shared; and that the anxiety created by the pace of change is fully taken on board and mitigated by the governance systems that we put in place.

Currently, the system is moving forward largely in silos. There are many different people talking about these topics in many different contexts, not sharing best practice and they are not learning from one another. Connection is important in a situation where data from one sector actually has huge value in a completely unrelated sector.

How do we navigate tensions created by capturing the benefits whilst distributing risks as evenly as possible, along with those equally distributed benefits? How do we improve public services whilst protecting privacy? These kinds of tension need to be out there in the open, openly discussed with very clear and accessible routes for people to understand how those discussions are being held. We need a safe space, in which the range of different stakeholders are deeply engaged in discussions and decisions.

There are challenges with the concepts that often underpin engagement. Ideas like ownership of data, privacy and consent are under incredible strain in the current system. What new concepts do we need to develop new frameworks that protect people in the way that they would like to feel protected, but at the same time support and capture the benefits of these technologies?

What is very clear is that across the sector there is a huge amount of uncertainty. There are some actors who are just racing forward regardless, whilst there are other actors who are kind of paralysed into not quite knowing what to do.

To shepherd this system through into the engaged, inclusive culture that we think is necessary, the Royal Society and British Academy proposed that there should be a new body, a stewardship body, which should be at the centre of thinking about this landscape. As a guiding principle for this body, the Academies argued that governance of data use should promote human flourishing. And layered under that are some really high level principles that should also inform all these debates. Protecting individuals but also thinking about collective rights and interests; having any trade offs discussed transparently, accountably and, crucially, inclusively; having systems in place to learn from the good practice of others but also to learn from bad practice and to correct errors or mistakes in the system; in general, to bring forward this whole concept of democratic governance that we hope will then be sufficiently flexible and robust to move forward as these technologies continue to develop.

What does this mean in practice? How are these principles to be applied in practice? What challenges should a stewardship body prioritize in navigating the data governance landscape and addressing the challenges presented above?

Data management and use: from principles to practice

The overarching principle of flourishing and the subsequent principles for data governance are likely to be applied differently across sectors. Understanding how these principles are currently applied and the lessons that can be drawn across sectors provides an initial glimpse of the whole governance landscape.

Health

There is huge potential for the use of health data for research and clinic practice. But mistakes have been made in the past, such as those that occurred with care.data, meaning that governance issues have been the focus of much discussion. Looking at examples from the health sector can help shed light on how to balance individual and collective rights in deciding how to use data for good.

Data governance in health is as much about enabling good things – uses of data that actively support flourishing – as it is about stopping bad things from happening. Research conducted by Wellcome and Understanding Patient Data shows that there tends to be support for uses of health data for research purposes where there was a perceived public benefit – though further research was under way to find out how people define public benefit in practice.

The principle of transparency is probably the biggest and most complex challenge for data governance in health. Transparency is important to have a sophisticated conversation, but there is no value in transparency if what is shared is not meaningful to people and there is generally very low awareness among the public of how data is used. However, Wellcome has found that, provided they were prompted to think about these issues using accessible and appropriate language, people actually wanted to know more. Explaining notions such as anonymisation and pseudonymisation to the general public, in a way that is honest and accurate and does not come across as patronising and disingenuous, is not trivial and demands attention.

Cities

City-level data can be valuable for addressing issues central to human flourishing. For example, using data to identify rogue landlord activities, can save money to the public purse and can improve a range of public health outcomes. Improving air quality in cities has an ostensibly positive, public good at the end of it which is to create a better understanding of air quality and how we measure it. Yet there are risks associated with data and its use. Measuring air quality brings in questions about the optimum configurations of IoT sensors, and some of the risks about that. Using data about passenger journeys based on access to wifi data raises issues of adequate consent.

In particular, with increasing use of algorithms, it is important to find ways to ensure they are both fair and accurate. Understanding the risks, benefits and end utility of algorithms might need serious and close-up user engagement as much as it would quantitative analysis of the data itself. Making trade-offs in complex ecosystems that are still subject to developing business models requires transparency and greater clarity around data-related responsibilities. To be able to put data to good use, cities would need in the first place to be very clear about how they would use the data and what they would use it for.

When it comes to the management of city data, what matters most is not creating a super data warehouse but rather to make sure that data from a range of sources is made available and usable, so that value can be extracted from it. Local authorities can also benefit from engaging with partner organisations to learn from and develop best practice, from academic institutions to business. With new opportunities come structures that public authorities need to consider more closely, such as data-sharing agreements, privacy-impact assessments, or cross-border sharing of data.

Services

The privacy paradox is presented as individuals feeling that their data is being abused whilst continuing to share their data with services that bring convenience. But this might be the most logical thing to do in a given situation – people judge that on balance signing up to a service can be beneficial, even if not entirely risk free. The overall situation could be much improved if organisations made decisions about data management and use more transparently and inclusively – giving people a chance to spell out what they deemed acceptable or unacceptable use of their data.

Better engagement requires more clarity and developing a common understanding about what is meant by terms like consent, personal data and misuse of data. The phrases, understood in the context of data protection, have almost no connection with what the person on the street would think was happening with their data. There is currently a gap between data protection and consumer protection. One can sign up for a service that promises not to share personal data, which could then use the data on the same terms to infer personal preferences and sell that information to a third party, so that when the consumer went to that third party that organisation could infer many characteristics about them.

Using data in services brings to light some specific approaches to issues such as consent. Trade-offs are often framed in terms of protection of personal or corporate rights to privacy or property, versus some more amorphous future public benefit to society. Sometimes this is the right way to frame the debate – and is the norm in the context of research. However, audit focuses on the safety of individuals in the here and now, currently utilising a service, requiring a different way of addressing tensions in governance of data.

National security

National security can have very different meanings in different countries. The primary purpose for the UK's national security agencies is to help keep the country safe. This requires data on people who are believed to be up to no good, so-called 'subjects of interest', and the purpose of collecting this data is to let society as a whole flourish – allowing people go about their everyday lives with the liberties and freedoms that they expect.

The trade-off between individual and collective rights is at the heart of national security. Necessity and proportionality are key concepts that are embedded into decision making, with actions that may be beneficial not being carried out, because they are not deemed to be proportionate. There are also considerations about consent that are particular to national security. It would be impossible for agencies to conduct an investigation if they had to ask the subjects of interest for consent to collect data about them, and here weighing up collective and individual rights is really fundamental.

Transparency and accountability is provided through a system of regulation and oversight. The Regulation of Investigatory Powers Act (RIPA), parliamentary and security committees, surveillance commissioners, and the Information Commissioner have a role in overseeing use of the data in the security sector and also in making sure that any actions taken are necessary and proportionate. But much of that process, for understandable reasons, is not in the public domain and the degree of transparency would always be somewhat limited.

There is a lot of learning from good practice going on in the national security space, which is now much more open than it used to be. There is now substantial engagement with a whole range of other government departments, with businesses, with academia about the way that agencies operated, and to develop a common understanding of what it meant for data to be held and to be used in appropriate ways. Dialogue on what is socially responsible and acceptable would be very valuable.

Data governance challenges: identifying and responding to the big issues

The establishment of the Human Fertilisation and Embryology Authority (HFEA) has proven highly effective in a debate with potential to provoke strong reactions, and that could have led to major backlash and poor outcomes. The HFEA has been successful in taking principles, steering a debate, bringing all on board and allowing progress in areas like in vitro fertilisation (IVF). Early action, as in this case, can be very effective and lead to sophisticated debates and significant societal progress.

As with the industrial revolution, governance lags behind the revolution brought about by data. The industrial revolution started in a fairly unregulated way but eventually there were a host of regulations that allowed industrialisation to continue and to prosper. The examples above show that we are now dealing with a much more complex landscape and potentially more profound transformations in society, in economies and other aspects. What are the key priorities for data governance in this context, and how can they bring about similarly successful governance?

Model leadership in safe and rapid innovation

Britain is a leader in both data-based innovation and responsible research. It has a thriving science and technology ecosystem. The Government has set out an ambition for the UK to be the best place to start up and scale and run a digital business and to be the safest and most secure place in all things digital. There is a strategic opportunity in the UK to make the country the most sensible place for companies to locate and start up a business, not because it is the least regulated but because it is the most sensible in the way it regulates.

The UK could develop its culture of ethics even further, by building deeper connections between data science and social sciences, and by asking hard questions. This included figuring out exactly what the country is trying to secure, and to have deep debates about evolving concepts such as privacy and autonomy.

Build on, and connect, existing governance

The starting point for governance is data protection and data protection law. The UK has a very strong legal framework that is currently being updated and we have a world-class regulator with strong regulatory powers. And, for the first time in a generation, the data protection legal regime is being modernised. The ICO will be able to take more proactive action than we ever have before. Under the GDPR (General Data Protection Regulation) there are codes of conduct and certification, mandatory and compulsory audits.

But there is a gap for a body that coordinates across multiple regulators. There are also gaps in the rules around the use of non-personal data, and the safeguarding of information that is derived from personal data. Once data has been anonymised, then the rules of data protection fall away, so there is a governance need concerning non-personal data. Critically, we need a stewardship body that works across all sectors and disciplines, including those that do not necessarily think of themselves as in the data sector.

Move from law to principles

Law always lags behind technology and law often lags behind social expectations and social norms. Data protection relies on solid principles of fairness, transparency and accountability. There is room to take those principles and apply them beyond personal data – to the use of data in smart cities, smart metering and autonomous vehicles has for communities at large. There is value in developing common codes of practice similar to those developed by professional bodies, councils or bodies like the Advertising Standards Authority, which is distinct from the regulator.

Practice foresight

We need to look across the landscape at the horizon: what is going on? What are the problems? What needs doing? There is a gap in foresight that is available to take a look at the ethical challenges ahead that draws information and expertise from a wide variety of disciplines.

Anticipating future governance requires a deliberative body to ensure foresight. The Nuffield Foundation's Convention proposes to look into such long and medium-term issues. This initiative would need to be independent from government, involve practitioners, and could be modelled on the Nuffield Council on Bioethics.

Foster dialogue and deliberation

Getting governance right in the 21st century, developing some new norms as a society and building public confidence requires inclusive, high-quality and authentic public engagement. A continued, two-way engagement would also allow decision makers to monitor public opinion of what is reasonable and fair, and to make decisions accordingly. Effective engagement is not trivial, and it could benefit from communicating both about good uses of data and about tangible harms that can arise from not using data.

Deliberation however goes beyond public engagement. Convening a space that connects the research of data scientists, the research of other disciplines – social scientists and philosophers – with those working in industry and in the public and private sector can complement public dialogue to build rich deliberation and debate.

Enable access to data

Some of the challenges of our time, including social and health care, could lead to the view that it is no longer acceptable that that data is not shared and exploited to give individuals and households the best possible chance of gaining the best care. Innovation too requires access to data, and the biggest problem is access to data for innovators, companies and small start-ups.

How does a small company try to innovate around transport systems or healthcare in a city? How do they negotiate with a city council in order to get access to that data? How does the city council provide that data in a way that companies can use it to innovate? Data trusts, presented in a Government commissioned AI review, could provide a framework in which companies are helped to negotiate with a city council, a big company or the Government departments. Emerging privacy preserving technologies can enable access to data while protecting sensitive information.

Tell good stories

There is a need to tell stories about the use of data and exemplar projects around good use of data because that helps the public gain confidence in how data can be used. A body that goes beyond regulation can provide a forum to tell such stories.

Stay agile

Crucially, in order to be able to deal with the unprecedented speed and nature of changes brought about by data and its use, any stewardship body would need to be formed with a flexible structure and retain the ability to adapt over time. As presented at the outset of the seminar, the world enabled by data is changing quickly and will continue to do so.

C: Seminar papers

The papers that follow have been produced in order to stimulate debate. They present various perspectives on what should be the priorities for a cross-sector governance initiative that would ensure the realisation of critical governance functions related to data management and use in the UK. This follows the recent publication of the British Academy and the Royal Society review, *Data management and use: governance in the 21st century*, which argues for the need to provide stewardship of the data governance landscape as a whole.

Several themes emerge from these contributions, including the following needs:

- Having a systemic approach, connecting debates further, and sharing learnings across the governance landscape.
- Identifying and addressing a limited number of issues to address in priority.
- Experimenting to determine what the operational model of an effective stewardship body should be.
- Building in transparency and engagement with all key stakeholders and the public, to ensure the stewardship body retains trust and legitimacy.
- Realising and sharing the value of data. This includes unlocking silos so that, for example, the most is made out of public data. It also includes exploring different models for ownership and consent.
- Providing an ethical framework, including for how organisations should develop and use data analytics such as machine learning and artificial intelligence.
- Reviewing the evolution of regulations and recommending new legislation.
- Demonstrating global leadership.

Michelle Brook and Anthony Zacharzewski

The Democratic Society

To rework the old joke about politics: you may not be interested in big data, but big data is interested in you. The new stewardship body has to take into account the views of the many actors interested in the current and emerging uses of data – academic researchers, industries, digital tool makers, NGOs and charities. It also must not forget the even greater numbers who are not that interested: the citizens.

It's not just in old jokes that data looks like politics. Like political power, data is widely dispersed, poorly explained and poorly understood. Its governance has to handle multiple conflicting and overlapping interests, in defence of an uninterested public.

There is a temptation to keep this an expert issue, to have good conversations with NGOs, academics, and charities, with commercial data users and innovative startups. The disagreement may be vivid, but it will always be well-informed.

We believe that this would be a serious mistake. If the stewardship body is to retain trust and legitimacy for its recommendations and actions, it needs to ensure that it hears from citizens – not just through campaign groups or representatives, but directly from people with a diverse range of experiences, insights, and values. It has to design its engagement methods – whatever its governance structure – for the age of 'show me' not the age of 'trust me'.

What does this look like?

The stewardship body has to aim for a way of working that brings diverse citizen voices into the conversation in different ways, and ensures that they are present and considered in every discussion. This starts from culture not structure. Creating an organisation which expects to involve citizens in decisions, and writes and thinks in accessible ways, will support the longer-term engagement and deeper conversations that the organisation will need.

Culture is repeated behaviour, and a sense of 'the way we do things here', so the first step is the right internal rules, rooted in the democratic principles of transparency, accountability, and enabling participation.

Learning from good practice that The Democratic Society, and many other organisations, recommend to governing institutions at local, national and international levels, core public statements and founding documents must ensure strong commitments to transparency and accountability. This is all the more important in policy areas that could drive fundamental shifts in our understandings of key concepts such as privacy and consent. Industrial actors have the potential to derive significant personal and commercial benefit – thorough transparency will enable engaged citizens and civil society to understand the trade offs and interests that are at stake, and the route through which decisions have been taken.

In practice, this means transparency of meetings (eg ensuring every significant decision is taken in the open, at a streamed meeting where possible, and with publication of minutes and supporting evidence). Powerful, visibly enforced codes of conduct on conflicts of interest must be in place from day one. Experts from the many differing domains of data practice and theory will inevitably bring personal and commercial views and conflicts of interest. The stewardship body should acknowledge these and work around them, rather than pretending they do not exist.

Transparency must also be effective, if it is to be meaningful. A thousand-page technical report is not transparent just because it is on a website. The audience for transparency is a non-technical one, and it requires a consistent effort to find a language and communications strategy that works for citizens, and to develop it to keep up with a rapidly-developing field.

Alongside transparency and accountability, the stewardship body must also design participation into its processes – starting with experimentation on what forms of participation work for this technical field. This is a challenging problem, that needs time to develop, but where the right culture and attitudes are essential. Use of lay members to represent citizen voice is unsatisfactory. It risks these individuals becoming sufficiently expert that they start to take on the arguments of one actors or another, and stop being a check and challenge. Additionally, a few citizens can never provide the diverse perspectives and insights that is one of the most valuable aspects of citizen participation.

Although the right structures need design and experimentation – an important short-term action – there are some key principles which are already clear. Participation should not be on a ‘who turns up’ approach but a more representative one, with reach into diverse communities of interest, demography and geography, and support for those who feel less confident to engage and be active. Formal mechanisms like citizen juries may be of use on some significant issues, but their expense rules them out in the day-to-day, so online panels or other quick and cheap methods of seeking representative public views could be used. Tools that reveal priorities and trade-offs will be particularly important, and will provide better feedback than an empty comment box. Framing of questions is essential, and needs to involve independent expertise to avoid unconscious bias. Citizen voice must be brought in throughout the policy process, from issue identification to implementation feedback, ‘before the beginning and after the end’.

Over the short and medium term the stewardship body has to clearly demonstrate that public views are being heard and having influence.

Simon Burall

Involve

This note, written at the request of the Royal Society, focuses on public engagement as one of the six key characteristics of the proposed stewardship body: deeply connected to diverse communities to create dialogue with and between publics, industry, civil society, academia and government.

Why is public engagement a key characteristic?

“At times...interests [regarding data management and use] will conflict, so openness is likely to be a key factor in enabling proximity to action while remaining truly and visibly free from capture by any one group.” (page 75 of *Data management and use: governance in the 21st century*)

The challenge for the stewardship body is how to understand and support the negotiation between different interests. This is made more difficult because our understanding of the current diversity of perspective and interest across multiple communities and domains of data use and management is poor. It is compounded by the fact that the perspectives of the multiple publics are dynamic and responsive to immediate events (both national and international) and over time, for example:

- different age groups appear to have different perspectives about privacy, but is this a cohort effect or the result of their stage of life?
- as data management and use embeds itself, society's reactions will change as the diverse nature of its impacts on different communities becomes more visible.

Effective social engagement by the data governance architecture should help reduce the strain which the report identifies on governance concepts such as privacy, ownership and consent. Given that the impacts of data use and management in different domains of data (health, education, defence etc) use will both affect publics' perspectives in the relevant domain as well as across the system, systemic stewardship is critical.

A stewardship body doing the public engagement

The risks

“...effective data governance strongly resists a one-size-fits-all approach, grounding efforts in underlying principles will provide a source of clarity and of trust across application areas.” (page 7 of *Data management and use: governance in the 21st century*)

As noted above, perspectives on data management and use will vary between communities, domains of use and over time. The bodies actually controlling and using the data will need a deep and instinctive understanding of the complexities of these views. To do this they need to remain connected to and able to understand and interpret changing public perspectives within their own contexts. They need to engage themselves and not sub-contract it to a central body.

The fact that society's reactions will be dynamic over time and that what happens in one domain will impact on and influence perspectives across the system, points to the role for the stewardship body being one of synthesising and building understanding of evolving public perspectives across domains, supporting those in individual domains to take account of the implications and cajoling domains and individual organisations to remain open and responsive to public perspectives.

What should the stewardship body actually do?

Short term: support the open and transparent development of a dynamic, shared framework for understanding and responding to public perspectives. This should include active co-production with key stakeholders across the domains of data management and use, and the public.

Medium term: support domains and individual organisations to build the capacity to engage with and become responsive to the public. Develop methodologies for synthesising understanding of publics' evolving perspectives in the system as a whole and support domains and individual organisations to respond effectively.

Long term: develop a longer-term picture of society's dynamic and changing perspectives. Governance architectures often become less effective over time as different centres of influence emerge. It will have a role in supporting the system to regularly step back and examine whether the system is still functioning according to the agreed principles.

About Involve

Involve was founded in 2003 to 'create a new focus for thinking and action on the links between new forms of public participation and existing democratic institutions'. Since then, we have developed world renowned expertise, skills and resources on public participation, open governance and democratic reform. We have partnered with international, national, devolved and local organisations – including the OECD, Open Government Partnership, UK and devolved governments and parliaments, numerous local authorities and public-sector bodies, and civil society organisations to:

- Develop innovative mechanisms for involving people in decision making;
- secure important reforms on government and corporate transparency and accountability;
- help build global norms and standards of openness and participation.

Estelle Clark

The Chartered Quality Institute

Introduction

I thank the British Academy / Royal Society for giving me the chance to offer my views about this vital subject and have relished the opportunity to briefly metamorphose into the chair of the new stewardship body!

I am highly impressed with the British Academy / Royal Society report, and associated documents, and understand the huge efforts that have gone into creating them. I also agree with the key assertion of the report; that human flourishing is the overarching principle that should guide the development of systems of governance.

It is not my place, nor could I, summarise all the considerations into this short paper. Rather I have distilled my thoughts into more operational concerns as the new stewardship body needs to be active as soon as possible and, recognising that I sound like the White Rabbit, is arguably already late!

Beliefs

Before I cover the short, medium and long-term I'd like to explain some overarching beliefs that inform my paper, and to consider the meaning of stewardship.

1. My starting point is that the world described in the report is one that needs to be inclusive. Inclusive in terms of everyone being able to be aware of, understand and control their lives and livelihoods in a world of new technology. And, inclusive in the sense of all strata of society being able to influence the debate and to inform the development of the new world ecosystem.
2. I fundamentally believe that the opportunities offered by the development of new technology, and the management of data, can be for the good of society; whilst also understanding that conspiracy, collusion or cock-up might cause the opposite to be true. Quite a conundrum, as the worst thing that we can do is to try to tame this new world. Rules can easily hamper those aiming to do good, whilst being easy to evade by those that desire to misbehave. And, we must remember that, there are no obvious national or international boundaries in this new world. Principles rather than rules must be the mantra and agile, adaptive systems our aim.

3. The UK is already a leader in the digital economy and I desire that the new body works to maximise the benefit to the UK wealth (in all its senses) recognising the UK's position not only as a technology leader but also in the province of good governance and the development of practical open standards.

4. Stewardship is the responsibility to shepherd and safeguard the valuables of others, in this instance by safeguarding the lives and livelihoods of people in relation to data management and use. To do so the new body must:

- be orientated towards a future well beyond the time when any immediate issue has been solved - this is not a one-time fix scenario;
- be capable of self-evaluation; frequent self-critical evaluations need to be built into the structure;
- manage complexity and recognise the speed of societal change by being distributed as widely as possible, with no function being performed by any part of the whole if it could be reasonably be done by a more peripheral part;
- be infinitely malleable yet extremely durable, being capable of constant, self-generated, modification of form or function without sacrificing its essential nature or embodied principle; and
- embrace diversity and change, attracting people and partners comfortable with such conditions and providing an environment in which they can flourish.

See table on next page

	Short term 0 – 6 months	Medium term 6 – 24 months	Long term 24+ months
Governance	<ul style="list-style-type: none"> Engage with UK population on ‘what good looks like’ for the new body Agree the principles by which this body exists Create the board and any advisory bodies Establish a long-term budget Scan the wider international set of institutions and standards 	<ul style="list-style-type: none"> Engage with UK population on ‘what good looks like’ for the new body Make decisions as to how the UK participates globally, including international agreement about the need for new standards 	<ul style="list-style-type: none"> Work with international bodies on the creation of new standards
Assurance	<ul style="list-style-type: none"> Start to build an organisation, based on the ideas in page 1, that the population can trust Consider how data management will be used in the operation of the stewardship body itself Define the metrics that stakeholders will use to assess performance of the development of the new body 	<ul style="list-style-type: none"> Complete build of the new organisation Use data management to the full in the stewardship body Use internal data to measure performance of the set-up of the new body Gauge opinions of full set of stakeholders on the set-up of the new body Define the metrics that stakeholders will use to assess performance of the operation of the new body Establish enforcement and remediation processes Openly report on progress 	<ul style="list-style-type: none"> Use data to measure performance of the new body against the original ‘what good looks like’ criteria Gauge opinions of full set of stakeholders as to whether data governance and use supports human flourishing Review enforcement and remediation processes Openly report on progress
Improvement	<ul style="list-style-type: none"> Take account of lessons learned from the establishment of other new systems of governance 	<ul style="list-style-type: none"> Review performance and feed ideas into rigorous improvement system Create case studies of good and bad practice in the establishment of data governance to aid learning in UK and elsewhere 	<ul style="list-style-type: none"> Review performance and feed ideas into rigorous improvement system Create case studies of good and bad practice in the establishment of data governance to aid learning in UK and elsewhere Review the efficiency and effectiveness of the improvement system
Leadership	<ul style="list-style-type: none"> Establish belief that the UK can take a leading role Establish the desire to involve the widest representation of society 	<ul style="list-style-type: none"> Ensure that the UK is taking a leading role Ensure that all strata of society are included without a single Cinderella 	<ul style="list-style-type: none"> Ensure the UK supports the development of other nations taking a leading role (complex, adaptive, system)
Context	<ul style="list-style-type: none"> Partner with existing UK data and technical governance bodies Understand the various sector interests Identify and link to technology leaders 	<ul style="list-style-type: none"> Partner with existing international data and technical governance bodies Identify and link to technology leaders 	<ul style="list-style-type: none"> Understand the various sector interests Identify and link to technology leaders

Guy Cohen

Privitar

Short term – Identifying, engaging and advising on emerging risks

New systemic risks to Critical National Infrastructure (CNI). As CNI is increasingly digitised, new risks emerge from the capacity for actions to be taken remotely, quickly and at scale. For example, with smart meters a home's power can be controlled remotely, creating the possibility for an attacker to cause a power surge by turning on many homes' devices at once. In financial markets consider the 2010 flash crash. The Body should engage with relevant authorities and develop a mechanism to report to those authorities on any emerging risks they encounter.

Direct threats to individuals. Big Data Analytics (BDA) can be unfair, opaque and harmful to citizen's privacy. The new stewardship body (Body) should support the ICO in its continuing work to protect individuals by developing a framework for how organisations should use BDA. The framework should include tests for fairness, a study of leading methods for increasing the transparency of opaque methods, and a comprehensive guide to privacy enhancing technologies (PETs).

Responding to changing liability. Digital disruption through AI can significantly change business models, at times this can mean that social functions once provided by these industries are no longer available and alternatives are required. For example, insurance provides a mechanism for individuals to share risk, which none could absorb independently. However, increased access to data and analytics will allow insurance companies to identify risk on an individual level to a much higher degree, meaning some individuals may become uninsurable and not have access to this protection. On the other side, growing numbers of connected devices through the IoT may require new liability models. Who is responsible for patching a second hand car and for how long? Who is liable for what when two automated vehicles crash? The Body should work to identify risks of this nature, and then engage with industry and Government to plan for solutions.

A new strategy to deal with antitrust issues. As industries digitise there is an increased propensity for monopolies to emerge due to a combination of network effects and benefits of scale. Consider law firms. Digitisation may increase the likelihood of a monopoly emerging. Partially this could be due to new network effects, for instance, disclosure is a time consuming process which requires many man hours as lawyers review legal documents to appraise relevance. As this process becomes automated, organisations which have seen more documents in various areas may be better at identifying what is relevant. Additionally, intelligent automated systems are cheaper at scale when considering capital costs of development; doubling the number of employees is around double the cost, doubling the application of an existing software product is not. The Body should review antitrust tools and explore alternative ways in which competition can be ensured and maintained.

Responding to emerging inequality. A combination of more winner takes most markets and automation threaten to further increase inequality on certain measures. BDA increases the share of economic benefit taken by capital and reduces that for labour. The Body should explore how this is happening and identify potential responses. These responses should include: evaluating the impact of increasing taxation on capital and reducing it on labour; requiring organisations to contribute to the retraining of the technological unemployed; changes to personal data rights which would aim to allow individuals to benefit from data they have provided; and distributed models which enable individuals to retain control and benefit from their data (eg Hub of All Things).

Restrictions on the use of BDA for specific purposes. There are certain applications of BDA which potentially should be restricted or banned. For instance, fully autonomous lethal weapons are potentially unethical due to the scale and speed with which they can kill. The Body should engage with relevant organisations to develop a list of applications of BDA which require Government approval, and then work with Government to review applications within those domains.

Medium Term – advocacy and education

Advocacy. Whilst real threats exist, negative and exaggerated headlines relating to improbable threats posed by AI may reduce the adoption of BDA and so slow the huge benefits BDA tools can deliver. The Body should act to mitigate this risk by advocating for BDA tools.

Education. BDA tools are likely to have an enormous impact on society, but can be challenging to understand and develop rapidly. There is a risk that a large portion of the population will be disenfranchised. This would both be harmful to those individuals, and, if that lack of understanding and engagement leads to alternative narratives, damaging to society more widely. The Body should act to prevent this by working to help the wider public understand how BDA tools work. Efforts should be made to work with all ages and communities, and to respond to misleading or inaccurate representations of BDA tools and projects.

Long term – the data social contract, information capital, and holding HMG to account.

Reviewing the data social contract. Historically citizens have contributed to society through taxes and obeying laws, and in return benefitted from the state services. Now, data offers a new way of contributing, with commensurate benefits. Most rights are not absolute; compulsory purchase orders are a limit on property rights, and imprisonment is a limit on the right to freedom. Privacy rights are still in their infancy, due to privacy threats from BDA being a relatively recent occurrence. What should the constraints on personal privacy be? Consider a rare and dangerous condition, should those who have been successfully treated retain their right to privacy, if their data could save the lives of others? These issues should be explored and unpacked by the Body with the intention of a legislative change with appropriate public debate and democratic accountability.

Exploring information capital. Data is increasingly referred to as the most valuable asset of today's world, yet it doesn't appear on company books. By failing to account for the value of data, we allow asymmetric exchanges of personal data for free services, we fail to realise the value of data and so don't protect and invest in data assets appropriately, and we limit the actions we can take to regulate and tax data rich organisations. There are clearly issues with alternative concepts of data ownership and existing models of IP and licensing, but the current constructs are also not fit for purpose. The Body should examine the issue, aiming to advise HMG on the advantages and disadvantages of the various models.

Holding HMG to account. In recent years HMG has significantly increased the extent to which it can process personal data (Better Use of Data/Digital Economy Act), with relatively little debate or public engagement. BDA tools allow for the repurposing of data which was initially granted for specific purposes in ways not imagined by the parliamentarians who approved the data collection initially. The re-evaluation of the social contract should go both ways, with HMG's use of data being thoroughly and continuously reviewed and their policies held to account. The Body should work with HMG to develop a framework for ethical use of public data and should be able to review existing HMG projects to ensure their data processing is proportionate and necessary. This framework should draw on the tools recommended by the ICO for use by private companies to safeguard personal data.

Matt Fenech

Future Advocacy

Introduction

The Royal Society and British Academy's report on Data management and use: Governance in the 21st century was a thoughtful addition to the important debate on the responsible use of data. By suggesting four principles of data governance, the conversation has been driven towards a more human-centric view of data governance, ensuring that all data use enables the 'promotion of human flourishing'¹. At Future Advocacy, we welcome the recommendations made in this report, including the creation of a new body 'to steward the evolution of the governance landscape'.

The need for such a body is made clearer when one considers the speed and scale of changes in this landscape. Over the last decade there's been an incredible increase in the data generated and retained by individuals, governments and companies. It has been estimated that 2.5 quintillion bytes of data are generated daily, and that more than 90% of the data in the world today has been created in the last four years². The business models of the internet giants of today, including Google, Facebook and Amazon, are dependent on the analysis and commercialisation of data provided by their users. There is also the potential for this data to be used in public policy-making and implementation, and in public service delivery. In its 2016 report commissioned by the European Commission, Deloitte identified 103 cases of big data analytics by public institutions worldwide³. Furthermore, these technologies have the potential to equip individuals and civil society with the necessary tools to participate in public decision-making processes, challenging the traditional hierarchies of public administration.

If you were chair of the new stewardship body, what would be your priorities?

Short-term priorities

The immediate priority for the community of academics, civil society organisations and other stakeholders in the data governance space concerns the implications of the changes in the UK legislative framework that are currently underway. Briefly, the EU's General Data Protection Regulation (GDPR) comes into effect in all EU member states on 25th May 2018. In parallel, the UK Government has published a draft Data Protection Bill (HL Bill 66), which is currently being debated in the House of Lords. When enacted in law, this Bill is intended to replace the Data Protection Act 1998, and also to transfer the provisions of the GDPR onto the UK Statute Book, in preparation for the UK's exit from the European Union in 2019. We have been made aware by our collaborators of a number of concerns surrounding the current draft of the Data Protection Bill. These include:

- A lack of explicit transparency safeguards in automated decision-making - that is, the so-called 'right to an explanation', protected by GDPR, is missing.
- The lack of an Article 80 derogation carried over from GDPR. This means that bodies cannot bring class action complaints to a regulator, unlike the situation in Germany, for example.
- The current draft appears to criminalise 'recklessly' re-identifying de-identified personal data without controller consent. Security researchers need further clarity that they will not be prosecuted for undertaking research that 'stress tests' the law.

1. The four principles suggested are to: protect individual and collective rights and interests; ensure that trade-offs affected by data management and data use are made transparently, accountably and inclusively; seek out good practices and learn from success and failure; and enhance existing democratic governance

2. IBM 'Bringing big data to enterprise', available at <https://www-01.ibm.com/software/data/bigdata/what-is-big-data.html>

3. Deloitte (2016) 'Big data analytics for policy making', available at https://joinup.ec.europa.eu/sites/default/files/dg_digit_study_big_data_analytics_for_policy_making.pdf

We believe that these concerns raise the possibility that a new Data Protection Act comes into conflict with the Royal Society and British Academy's principles for good data governance, particularly the principle that we should 'ensure that trade-offs affected by data management and data use are made transparently, accountably and inclusively'. Although a stewardship body may not wish to undertake a direct advocacy/lobbying function, it may have a role in acting as a convener for the various stakeholders with an interest in ensuring that UK legislation is fit for purpose. It may bring together diverse views, ensure the distillation of a clear message, and act as a powerful amplifier for this message. In this way, it will ensure that the principles outlined in the Royal Society and British Academy's report are upheld in law.

Medium- to long-term priorities

The increasing collection and use of data has allowed incredible developments in both the private and public sectors, as outlined above. Such developments are not without risk, however. From data loss by companies and public bodies (eg Suffolk County Council in 2013⁴, Yahoo in 2013⁵, TalkTalk in 2015⁶) with attendant risk of identity theft and fraudulent financial activity, to the potential for surveillance and espionage activity by governments on their own citizens (as outlined by the Snowden revelations in 2013⁷), new data analytics techniques can allow information about groups of people and indeed specific individuals to be gleaned without the person's express consent, or indeed knowledge. There are also the perhaps more pragmatic issues to be considered:

1. To access so many services/amenities in the modern world, consumers need to supply large amounts of their personal data, very often in perpetuity and with little control or knowledge of how this data is used, shared and stored. There is a tension between this 'new normal' in our online interactions and a widely-shared need for privacy and autonomy when it comes to our sense of identity.
2. The discussion around data provision by consumers and use by private companies has become much more explicitly commercial over the last few years, but few users reflect on the implications of this and whether they are indeed getting value out of these data transactions. Put another way: are current data transactions between consumers/citizens and private/public bodies fair, with all parties extracting equal or comparable value from the interaction?

4. <http://www.ipswichstar.co.uk/news/ipswich-suffolk-inquiry-starts-after-memory-stick-lost-by-county-1-3058290>

5. <http://www.telegraph.co.uk/technology/2016/12/15/yahoo-hack-need-know-biggest-data-breach-history/>

6. <https://www.theguardian.com/business/2015/nov/06/nearly-157000-had-data-breached-in-talktalk-cyber-attack>

7. Lyon D (2014) 'Surveillance, Snowden, and Big Data: Capacities, consequences, critique', *Big Data and Society*, available at <http://journals.sagepub.com/doi/abs/10.1177/2053951714541861>

A stewardship body on Data Governance should seek to address these questions in the medium- to long-term. It should work to ensure that data transactions between consumers and companies operate as a partnership. This partnership will entail a more equal relationship between participants in data transactions, where citizens are clear about:

1. How their data are being used (how/where their data is stored, how long for, who it is shared with) by companies and governments; and
2. What benefits they can expect to receive in return for the huge commercial benefits companies receive from access to this data. These benefits can be collectively termed the 'digital dividend'.

In the ideal world that the stewardship body will be working towards, each citizen will be well-informed and empowered enough to take control of their own data and ensure they are deriving appropriate benefit from it. In short, the stewardship body will recommend courses of action that will guarantee that data collection and use benefits everyone equally - this is the central principle that underpins what we have termed a 'New Deal on Data'⁸.

About Future Advocacy

Future Advocacy is a think tank focused on making sure the United Kingdom is best positioned to capitalize on the opportunities and mitigate the risks presented by artificial intelligence, big data, and similar technologies driving the Intelligence Revolution. Our vision is a world in which the social, ethical and economic opportunities of AI and big data are maximised, while the risks are minimised. Our mission is to work collaboratively to advocate successfully for the policy changes, business practice changes, and individual behavioural changes that will ensure that AI development, and the use of data underpinning this, is beneficial to all humanity.

8. Future Advocacy (2016) 'An Intelligent Future? Maximising the opportunities and minimising the risks of artificial intelligence in the UK', available at <https://www.futureadvocacy.com/s/An-intelligent-future-3.pdf>

Tim Hill

Law Society

The core functions for a successful governance landscape could shape the short, medium and long-term priorities of the new stewardship body.

In 1942 the Beveridge Report anticipated the task of post-war reconstruction in seeking to overcome five giants: squalor, ignorance, want, idleness, and disease. Its vision captured the public imagination; its practical implementation, the development of a socio-technical data processing system that promoted human flourishing and influenced the world: the welfare state.

Human flourishing in the 21st century, like the post-war task of reconstruction over 70 years ago, is still held back by the old giants, but we need no longer rely on the same old hierarchic and technocratic systems to fight them. Nor need we imagine that we can look solely to the market.

The first priorities of the new body should therefore be threefold:

1. anticipate what is becoming possible in the use of data to tackle social ills and present this as an achievable vision;
2. orchestrate democratic and inclusive debate to identify the actual tasks we wish to tackle as a society; and
3. challenge public and private actors and individuals to participate in making a reality of this vision.

The actual tasks could perhaps model themselves on the Grand Challenges – ambitious goals to be achieved within a certain timescale – using individual, corporate and public data to tackle a particular disease, an area of poverty, or an environmental problem.

The difficulty that will be encountered in addressing these challenges should segue into medium-term priority of building practices and standards: the particular practical governance frameworks – data looms that weave the web of personal data with the warp of non-personal data – that can deliver real benefits.

It is difficult to say now what these frameworks might look like. They will have been forged in practical fashion across a number of disciplines and with varying success. At this stage the priorities will be:

1. learning lessons from real world delivery;
2. generalisation of framework models; and
3. diffusion of workable governance models for adoption by others in tackling new challenges.

Having set a vision, learned lessons in seeking to making it a reality and promoted models for wider adoption (and as templates for new models), the new body might then turn fully to its core function of clarifying, enforcing and remedying.

Jackie Hunter

BenevolentBio

Situation

A review of the data landscape across all industries will identify a diverse range of types, formats and challenges associated with access that can impact our ability to make informed decisions. In the biomedical space the five highest value types of data that can be content mined and lead to novel insights and innovation are: research data (papers, abstracts, supplementary material), raw clinical data, raw biological data, patents and chemical structures. Within a UK context, research institutions and the NHS, are primarily or totally funded by government and hence the tax payer. The access to this information for innovation is extremely limited, especially for SMEs and therefore its ability to drive and fuel innovation in the UK is stifled. The costs of licensing content for example from publishers of this publically funded work is extremely high, especially for SMEs. It is well recognised that the current system of green and gold access for research publishers is not working (eg Green et al 2017 doi 10.1002/leap.1116) and the default stance, often enforced by publishers, is to deny access to the publically funded research for data mining purposes.

Clinical information clearly needs to maintain patient confidentiality and even though some data is made accessible in summary form, rarely is the raw data made available. For biological data the onus is on grant holders and institutions to maintain data repositories or make raw data available and there is little or no incentive from funders or institutions to do this eg lack of recognition in both RAE/REF and promotion purposes.

This situation can and should change and a data stewardship body could be an excellent way to facilitate this – done wrong and a bureaucratic top down body could completely stifle innovation.

Target

As a minimum to have all publicly funded data and information available for mining (via bulk download access, as opposed to publishers offering the ability to mine data, but using their infrastructure, platforms, and algorithms) and knowledge creation to drive value for the UK economy and increased health and prosperity of UK citizens.

Proposal

Short term

Of course a governing body of this nature needs to be independent and represent the needs of all stakeholders. In addition such a governing body needs to understand the issues and subtleties around data sharing and therefore it is important that this body has some experience of reduction to practice. It is desirable therefore that this body has some operational function as well as governance oversight. In order to gain some of this experience, the body needs to identify and define the key stakeholders and then first assess what data is available for mining and how this is being done (or not) across the different types of data mentioned above.

Data that is funded by taxpayers but is not available should also be identified.

Medium term

Put in place mechanisms and incentives to permit access to key data sources eg publications, clinical trial data, biological data repositories. This might be through a master deal with publishers as is happening in Germany, through the use of a 'doi' system to recognise data and data sharing which would count towards research assessment, making use of the funding mechanisms to drive change (a good example of this is Sally Davies's making Athena Swan accreditation for medical schools a condition of receiving NHR funding) and using publishers to demand that raw data is made available for mining for publications (as was the case with microarray data).

Establish the governing body as an honest broker in terms of operationalising data sharing initiatives.

Irene Ng

HAT Foundation Group

I write in response to your request for a provocation paper and attach below my views on the matter. My interest and experience in the personal data economy is on its economic impact, the design of the market and the alignment of incentives, as well as the engineering and design of personal data platforms, the data schema and structures. My experience has been to use technology, regulatory, market and civil society levers wherever they may be relevant to achieve the human flourishing objective as part of the mission of the HAT Community Foundation (<https://hatcommunity.org>). Of course, I am mindful that while choice and empowerment of the individual is key, it is for the individual to decide if he wants to flourish, instead of the decision taken for him by another entity, in the name of human flourishing. There is also consideration of the commons, and the necessary trade offs of commons v individual human flourishing.

As a foundational premise, I wish to acknowledge 2 separate types of personal data that require different approaches to stewardship.

Corporation held personal data (CPD)

This is 99.9% of personal data held currently. The practices are diverse. Some corporations do not collect the data (eg Telco) while some not only collect it, but run advanced analytics on it. Some SMEs are in fear of going anywhere near personal data as they feel they can't risk fallout. Many are shutting down login functionalities of websites. This potentially hands the market to businesses that operate outside the EU, through outsourcing contracts. On the other hand, other SMEs are hoovering up data to sell. The stewardship body could provide guidance and skills on data management and governance but also of innovation and opportunities.

Individually controlled personal data (IPD)

A fledgling group of private data accounts such as mydex, citizen.me, Digi.me, people.io, cozy.io, meeco.me as well as HAT (hubofallthings.com) private data accounts such as savy.io, nogginpod.me are bringing a wave of Individually controlled personal data. For example, Facebook data held by Facebook is CPD; the same Facebook data held by private data accounts are IPD. Individually controlled personal data are fully controlled by the individual; can generate new data through personal AI and are potentially a powerful force within the personal data economy. If data controller status is given to the individual (not all private data accounts are technologically or legally designed to do so), consultations with the ICO suggest that they may be also exempt from GDPR 2018.

Based on CPD and IPD, I would argue that personal data has 'polarity'. This means that for the SAME data, where it sits, how and where it's used, and by whom, will all have different value and different risks. There needs to be a recognition of that in data science and policy, which implies that approaches to data management and use is a combination of social science and science methodologies. The collaboration of Royal Society with the British Academy in conducting this review, a reflection of such awareness, is most commendable.

My thoughts on actions:

The Internet is thriving on the trade and exchange of personal data, legal or otherwise. Personal data fuels a €272 billion economy of ad and ad blockers, real and fake news, real and satirical media; In short, almost all activity online. Where there are illegal practices, the law can barely be enforced. An Internet service that finds it hard to comply with some practices can move itself to a different jurisdiction and still provide the same service. For a stewardship body to have real influence it needs to understand all the different levers to have a chance at being effective and it must also understand the limits of legislation and the potential moral hazard it creates for digital services. If compliance or stewardship guidance is costly both in terms of economic costs and goodwill from the market, digital services can find another jurisdiction, leaving only the organisations that comply to carry the cost of compliance while others will cream off revenues to a jurisdiction outside the boundaries of legislation. Not only will this reduce the competitiveness of national industries and organisations in the digital economy, but the state will also lose the ability to tax digital services. The recent cases of Amazon and Apple reinforce this.

For this reason, this provocation deals almost entirely with economic levers.

Short term

In the short term, it is important to consider treatment only of data that hold the highest risks, of which they are often also providing the greatest gains. In other words, it is not merely the content of the data that is risky eg personal data; but where they are held and how they are accessible. Highly personal data locked down without any access by the organisation outside of their own offices (ie not cloud enabled) would have lower risk than data that can be accessed. In other words, there is a need to address data that has high mobility – because mobility of data brings opportunities and risks. By self-selection, firms that allow their data to be most mobile would want to reap the greatest opportunities - they must be aware of the costs, and data governance must be aware of the economic levers for both risks and opportunities brought about by data that has high mobility.

The short-term priority is to engage, document and report the economic levers.

Engage with supply: Engage with tech and non-tech companies that hold personal data. Understand their costs, risks, and opportunities. Engage with IoT device makers and apps that are generating petabytes personal data, some anonymising and selling them to fund their activities. Others are expunging personal data to mitigate risks. Understand practices and motivations. Summarise the economic levers that would influence supply and potential interventions.

Engage with demand: Get a full understanding of what models are fuelling the 31 billion dollar Advertising tech market. Specifically, this market buys 'signals' which are transformed personal data that tells the marketer the person's propensity to receive what kind of messages, content, ads etc. Transformation of data to signals is what fuels the personal data economy. Raw personal data is still valuable, but its proportion of real worth is likely to be the same worth of a lump of charcoal to the beautiful diamond sold in retail shops. The multiplier effect from the multiple transformation levels (whether within one organisation such as Google, or in the market such as other big data or analytics firms selling insights) cannot be underestimated. Google profits from personal data not because they trade it, but because they can convert data into signals. Hence USD31b advertising economy can generate create a company with a market capitalisation of USD600b. A good understanding of the monetisation, revenue streams, and intermediaries is needed – see how complex the landscape is. <http://chiefmartec.com/2016/03/marketing-technology-landscape-supergraphic-2016/>

Summarise the economic levers that would influence demand and consider potential interventions

Engage with use

Many consumers do not even know that the underlying fuel for their digital services is personal data. The old saying of Uber having no cars, Airbnb no property and Facebook having no content just means that these services that are able to coordinate and use personal data or information about user contexts are an asset class of its own. An understanding of what personal data fuels what kind of use would be needed. Understanding the many perceptual and actual challenges around trust, identity, privacy, security, vulnerability, control is key. Summarise behavioural challenges and levers such as scripting, nudges and other design levers that can provide assurances, which alleviate perceived risk but on the other hand, could also heighten fear, create concerns and stop usage. Summarise potential interventions.

Engage with management practices, processes and policies: Consolidate some of the useful frameworks that have been developed by consultants and organisations and create a set of best practices.

Engage with tech giants on how they have been governing personal data to document practices.

Summarise best practices and management levers that impact supply, demand and use.

Short term, there need to be an empirical understanding of the personal data economy (PDE) and the personal data asset class. As your report has indicated, governance and management cannot be separated. I would argue that with the intense connectivity of both online and offline human personas, all personal data flows are not separable. The key to the stewardship body having any relevance or success is to have a systemic view, with systemic methodologies and to understand the levers for interventions that can assist to achieve the objectives.

Medium to long term

For the personal data economy model, the medium term is a richer and more robust model with data, able to provide better guidance, stewardship of personal data supply, demand, use and management for governments, policy makers, businesses, citizens.

It would be wise to begin empirical economic cost-benefit modelling (and other modelling initiatives) to stratify personal data to provide useful guidance on what personal data, sat in what environment, would be risky in what context, and with systemic model assumptions. The parameters are known, even in the IoT space⁹. Build scenario models - specifically societal impact, GDP impact.

Support the proliferation of Individually controlled Personal Data accounts – they provide genuine alternatives to the current Internet model of Corporate controlled Personal Data.

Work with consumer advocacy groups such as Which! (Who are starting to develop policies in the personal data area) and the ICO on ways to create a more transparent environment of data usage so that enforcement of wrong doing can be possible - both through market mechanisms such as consumer class actions, as well through penalties.

The HAT Foundation Group has provided guidance globally on data governance, in particular the economic models of personal data. It is independent and members owned, but also have oversight and guardianship over an open sourced technology platform for personal data exchange adopted and operated by several providers globally. For almost 6 years, we have been working on the personal data economy, and some of our briefing reports can be downloaded here: <https://hatresearch.org/hatoutputs/hat-briefing-papers/>

The foundation's agenda is fully aligned with the human flourishing agenda of the report. Indeed, the empowerment of individuals is central to the foundation's mission. With a full innovation programme of startups building on private data accounts, and a host of partners globally, we stand ready to contribute to the stewardship body.

9. Operationalising IoT for reverse supply: the development of use-visibility measures
<http://www.emeraldinsight.com/doi/full/10.1108/SCM-10-2015-0386>

Nora Ni Loideain

Information Law and Policy Centre at the Institute for Advanced Legal Studies

Greater power should come with greater accountability: revisiting the roles of data controllers and processors

The processing of personal data is becoming increasingly more complex and encompasses more tasks by different actors than ever before. In addition to the management and use of data by the public and private sectors separately, the rapidly growing ‘third space’ (where both sectors are involved at different key stages of the data processing cycle) is now providing for ever more entrenched data-driven systems that will increasingly underpin the governance of cities, homes, and essential public services, such as health, transport, and law enforcement.¹⁰

EU data protection law has traditionally (and continues to do so under the GDPR) accorded different levels of responsibility to the actors involved in this management and use of information who fall within the scope of ‘data controllers’ and ‘data processors’. The former concept comprises the party who alone or jointly with others ‘determines the purposes and means’ of the processing, whereas the latter is the actor that processes the data ‘on behalf’ of the data controller. In line with the requirements to implement the GDPR, this approach has also been adopted in the UK Draft Data Protection Bill 2017.

In light of Big Data processing, which is becoming more automated and reliant on data-driven decision-making (eg via the use of AI), and the trend of outsourcing any (especially complex) data processing task, the distinction between the roles and responsibilities of data controllers and processors has become less clear¹¹. In particular, the question of who is determining the purpose(s) of the processing in question within this particular context may be less than obvious.

Consequently, the relevant actor responsible for ensuring a culture of good data governance compliance (the data controller) may not be in a position to exercise the adequate and necessary levels of scrutiny, oversight, and enforcement over a data processor (eg a much more technically expert vendor)¹². Such opaque data processing arrangements invariably raise concerns for the protection of individual’s and society’s rights and interests and the implications of the ‘trade-offs’ resulting from these systems for the data protection principles of transparency and accountability, and good governance.

In order to ensure adequate and effective accountability across the data processing landscape, the concept of data controllers being the main actor tasked with the overall responsibility of ensuring compliance with good data governance should be reconsidered. For instance, if principles such as data protection by design and default (or privacy by design) are to be implemented at every stage of the data processing cycle, then each relevant actor should be directly legally required to implement technical and organisational measures, eg such as data minimisation.¹³

In other words, it is an insufficient guarantee of ensuring good data governance in the 21st century to simply state that producers of products, services and applications should be ‘encouraged’ to take into account the right to data protection when developing and designing such products, services and applications.

10. See e.g. Kitchin, R. (2016) Getting smarter about smart cities: Improving data privacy and data security. Data Protection Unit, Department of the Taoiseach, Dublin, Ireland.

11. See e.g. ICO, Big Data, AI, Machine Learning and Data Protection (Version 2.2., 2017), 56.

12. See e.g. N. Ni Loideain, ‘Cape Town as a Smart and Safe City: Implications for Governance and Data Privacy’ (2017) International Data Privacy Law (forthcoming).

13. This proposal departs from GDPR, art.25 which limits the implementation of data protection by design and default to the responsibility of data controllers.

Kieron O'Hara

University of Southampton

Privacy discourses and a Data Governance Stewardship Council

Data Management and Use, a joint report from the British Academy and the Royal Society, has recommended that a data governance stewardship council be set up to influence the evolution of the landscape, as data play an ever more prominent role in our daily lives. The report leaves open the question of what principles the council should be endeavouring to preserve, and what its priorities should be.

There are many potential functions that such a body could undertake. The report itself focuses on three broad groups of function: anticipation and evaluation of trends; building practices and setting standards; enforcing norms and remedying harms. In this short paper I don't want to enter directly into the debate about precisely where such a body should focus its resources and attention. Rather, I want to propose a framework for separating out and distinguishing the functions it might take on.

I will apply a framework currently under development to understand and disentangle the many overlapping discourses on privacy¹⁴, whose aim is to resolve at least some of the confusion, complexity and category error characteristic of privacy discussions¹⁵. Privacy is of course one of the central factors in data governance, indeed a factor which, according to the report, is "under unprecedented strain" (p2). I will therefore focus on privacy; however, the framework may be extendable to other aspects of data governance. Privacy is only a part of such a council's competence, but an important part.

The framework separates out seven distinct types of privacy discourse, which are often confused and inappropriately mixed. For each of these levels, we can ask should the data governance stewardship council operate at this level?

Level 1: concepts

Privacy takes many forms, and the various conceptions of what privacy consists in are the topic of level 1. In the realms of data, conceptions include informational privacy (Bob does not process data about Alice), decisional privacy (Bob doesn't interfere with Alice's decisions, eg by providing her with partial data), and economic privacy (Bob doesn't appropriate Alice's intellectual property). Discussion of privacy at level 1 concerns the nature of privacy relations implicated by particular uses of data.

There are few pressing functions at this level for the council. However, this would cover horizon scanning, to explore the privacy implications of alternative and possible futures and 'black swan' events, and also scrutinising the status quo to determine what privacy issues exist today that perhaps have not been surfaced.

Level 2: empirical measurement

Given a particular conception of privacy, the level 2 discussion is whether people in actual fact have privacy in that sense. Eg if the privacy notion under discussion is that Bob doesn't process Alice's data, then the level 2 discussion would be the fact of the matter of whether Bob actually does or does not process Alice's data. If the former, then Alice does not have privacy in this sense; if the latter, then she does. She may not care if she does not have privacy in this sense, or she may have provided her consent – these are irrelevant to the question of whether or not she actually has privacy.

If the council had level 2 functions, then these would connect with an ombudsman-like role, exercising quasi-legal powers to: demand information about their data processing from data controllers; conduct assessments and spot inspections; determine whether privacy policies were being followed; or mandate particular cybersecurity techniques.

14. Kieron O'Hara, 'The seven veils of privacy', IEEE Internet Computing, 20(2), 2016, 86-91.

15. Cf. e.g. Daniel Solove, Understanding Privacy, Harvard University Press, 2008.

Level 3: phenomenology

Level 3 contains the (relatively rare) discussions about what privacy (or its lack) actually feels like to an individual. Face to face privacy breaches tend to produce feelings of shame or embarrassment, but data breaches can be unnoticed. Facebook spends a lot of resource in making its users feel private even while they haemorrhage data in the general direction of Zuckerberg. Too-specific, too-timely adverts can give people feelings of creepiness. Users of smartphones feel that their phones are extensions of themselves, not that they are gathering and disseminating data. On the other hand, in many cases a lack of privacy can feel invigorating or exciting.

Level 3 functions for a council would be relatively rare, and unlikely to be its sole focus. However, there are numerous opportunities for Human-Computer Interaction research to change the status quo around data consent, but fundamental to such change is that people see a need for it. One may have very well defined terms and conditions, but if people don't even know their data is being accessed – if this use is not apparent in the first place – transparency about the process will add very little. Potential level 3 roles for the council would be to monitor data-users' advertising to ensure it is not misleading, or to promote designs which signal the fact of data use in some way.

Level 4: preferences

Level 4 contains discussion of individuals' privacy preferences. These can be entirely idiosyncratic, and need not be consistent, rational or reasonable. It is hard to generalise; sometimes people wish to be concealed, and sometimes they wish to be visible to their network. On other occasions, people might want privacy in the abstract, but have a greater preference for some other type of good (eg free services). At this level, we find ideas such as consent, personal data stores and privacy markets, which allow individuals to impose their own conditions on the use of their data.

If the council was to be focused around level 4, then we would expect its role to be to ensure that individuals' preferences about the use of 'their' data (however that phrase is understood) are respected.

Level 5: norms

At level 5, we find discussions about social norms of privacy, conventions, regularities, expectations. These norms are unenforced. To complicate the issue, they also vary across culture, generation and gender.

A focus at this level for the council would mean that its role was to preserve the reasonable expectations of privacy of citizens, and to help manage the contextual integrity or appropriacy of information flow.

Level 6: law and regulation

Level 6 is the level at which we discuss the legal questions of privacy – what privacy breaches are against the law, and when does the state mandate transparency? Jurisdictional issues are relevant here as well. Data protection (unlike privacy) is an entirely legal concept, and lives at this level.

This is perhaps the least likely level at which the council might find a role, as it would usurp functions of the courts (for privacy and rights judgements) and the Information Commissioner's Office (for data protection issues).

Level 7: politics and morality

The final level is that of the value of privacy, and when privacy is right or wrong. Issues here include whether privacy is essential for vital civic functions such as democratic deliberation, or psychological well-being and the development of individual autonomy, and whether privacy is damaging for other important social goods, such as security, or investigations into criminality or terrorism.

These are essentially political matters, so if the council worked at this level it would share an arena with Parliamentary politics and the media. However, it could help pronounce on, or shape opinion about, complex questions in this zone such as when concerns about privacy are trumped by the social value of, say, data-driven medical research.

The council could work at any or all of these levels, as deemed appropriate. However, it would also be reasonable to expect that it would work best if it was targeted directly at one or two levels, in order to ensure clear lines of control and power, and to establish its particular expertise. This paper is agnostic about which levels they may be. However, it is hoped that the 7-level framework outlined is helpful in framing the future discussion.

Hetan Shah

The Royal Statistical Society

Practical matters

The RS/BA report helpfully identifies the issues that need grappling with and some missing functions that need to be picked up. My view is that the more regulatory type issues should be housed in the Information Commissioner's Office rather than in an entirely new body. It would save time and cost, and would strengthen the role of the ICO at a time when it is much needed.

The more reflective functions could be picked up by the new Nuffield Foundation Convention on Data Ethics. The two bodies might develop a supportive working relationship similar to that of the Human Fertilisation and Embryology Authority and the Nuffield Council on Bioethics.

Priorities

On the regulatory side, the ICO type body might consider:

- Regulation of biometrics, in particular facial and voice recognition.
- What does GDPR miss and how should these gaps be dealt with? Eg concerns about inference which do not break DP rules.
- Making regulation across different data fields more coherent.
- Making the case for simplifying researcher access to administrative data.
- Breaking down silos for data sharing in government (particularly around health data).
- Consideration of what a regulator can expect in terms of transparency / explainability around algorithmic outcomes.
- Building capacity and knowledge amongst other regulators in specific industries of what new issues might arise through new data technologies such as machine learning / algorithms.

The Convention on Data Ethics should focus on a set of rolling questions (ie pick a few areas of focus, work on them and then move on). It might consider:

- Clarifying terms (as the Warnock commission did).
- Considering the philosophical question of what new ethical and governance issues arise from the changing nature of data, and which are the same old questions as before. From this to develop an ethical framework.
- Code of conduct for data science community (working in partnership with professional bodies such as the Royal Statistical Society).
- More deliberative dialogue consultations with the public particularly around data linkages (ie not just how they feel about any particular personal data being used, but how they feel when it is linked up).
- New models for regulating data monopolies – eg data as commons.
- Creating a safe space for, and providing guidance to private sector organisations on the ethics of particular potential projects.
- Bringing together those working on addressing bias in algorithms to develop guidance on how this can be done.
- Tackling specific issues as they appear (often tied to new technologies) – eg autonomous vehicles.
- Developing links with other bodies internationally that are working in this area.

If there are other matters outside of the regulatory and the deliberative which need picking up, one might consider if these should simply be tendered out as discrete projects, rather than incurring the costs of having to set up an entirely new body. There are many organisations working in this area (eg the Open Data Institute) which one could imagine might pick up any remaining issues if they were tendered with appropriate funding.

Amanda Smith, Jeni Tennison, Leigh Dodds and Peter Wells

Open Data Institute

The Royal Society has conducted a review on data governance for the UK. The review used a broad definition of data governance: “everything designed to inform the extent of confidence in data management, data use and the technologies derived from it”.

The report recommended the development of a set of high-level principles and creation of a stewardship body for the UK’s data governance landscape. They have asked participants at a roundtable to submit short, 1 – 2 page, responses to the following questions:

“If you were Chair of the new stewardship body, what would be your priorities, 1) short-term? 2) medium-term? 3) long-term?”

The Open Data Institute is a global organisation, headquartered in London. Our response is intended to supplement the Royal Society’s report and its recommendations for the data stewardship body.

Our response is also based on our belief that data is a tool for innovation and an emerging virtual and vital infrastructure that all countries and all sectors of the economy rely on. The data in this infrastructure should be as open as possible. Better access to data maximises the use and value of data but this must be balanced against controls to protect the privacy, interests and security of people, organisations and states.

Short-term (100 days)

1. Deepen understanding of the landscape

For the data steward to be effective it will need a deeper understanding of this organisational landscape and strong engagement with people and organisations across the country.

The review assessed the UK’s regulatory landscape but did not review the existing organisations that are already performing some of these data governance functions. There are many of these organisations both in particular sectors and at local, national and global levels. For example devolved administrations, local authorities and company/university internal ethics boards; national organisations for sports, banking and the Information Commissioner’s Office; or international organisations for academia and agriculture. UK government has agreed to create a new data ethics council.

Decisions about data should often be made by those closest to the decision: such as the people identifiable in or impacted by the data, the organisations working in those contexts or the appropriate democratically accountable public body. These organisations have existing capabilities, tools and practices. Their good practices could be extended into more contexts.

2. Start building peer networks

There are many ways that the data steward could choose to work with this landscape of organisations. We recommend building peer networks.

At the Open Data Institute we work with a peer network of ODI Nodes across the world. This includes nodes in all four UK nations and many UK city-regions. Each of the nodes help local and regional government, businesses, community groups, academia and innovators and provide a neutral space to convene different groups to solve data-related problems and grow a strong, fair and sustainable data economy in their locality. All of the nodes in the network share lessons with each other. We have similarly built peer networks across particular roles (Open Data Leaders in the public sector in multiple countries) and sectors (for example in sports and retail).

We recommend that the data steward follows a similar model and state that it will build a range of peer networks across the UK and across every sector of the economy. Our method report on building one of those networks may be useful.

3. Pick a specific challenge (or two, or three) to tackle first

There are many challenges identified in the report, and many challenges faced by the current data governance organisations. It is not possible to tackle them all immediately. There will need to be some focus along with the general direction of promoting human flourishing. That early focus will help establish the organisation and demonstrate its capability to create change. An early priority for the data stewardship should be to openly identify and clearly communicate the challenges it will initially focus on, how it intends to start tackling them, and how people and organisations can get engaged with the work.

One initial challenge may be to develop better approaches for certification by looking for learnings from and common patterns across the multiple approaches currently in place, engaging stakeholders to understand needs, and collaboratively developing a better process for the many new certification needs in coming years.

Another initial challenge may be devolution. The data stewardship body could lead research on where, why and which functions of data governance could be devolved, models in use in other countries, and suggest options for future devolution deals within the UK. Greater devolution for some data governance powers may allow greater democratic involvement through devolved national administrations and city-region mayors. This could help those regions to achieve their goals such as improving transport, reducing homelessness or increasing economic growth and jobs.

Medium-term (in 2 years)

1. Demonstrate value by tackling those challenges

By the end of its second year the data steward should be able to demonstrate progress on some or all of those initial challenges. It should have experimented to discover how it can create change and what its operating model should be. It should have shared learnings and case studies. It should have worked in the open to show the value of open culture in widening debate and encouraging more organisations to work together to solve common problems. It should have improved the lives of citizens.

2. Build the capacity of the networks

The peer networks of UK data governance organisations will be an asset. The data steward should understand their needs and build their capacity including, where necessary, supporting new organisations to emerge in particular geographic or sectoral contexts. It might need to source funding for new organisations for an initial period.

The data steward organisation should help disseminate learnings - both of what works, and what doesn't - across the network. This might take the form of training materials, workshops, case studies, etcetera.

3. Engage with citizens and understand their expectations

The review looked at existing studies on UK citizen's views on data use. It will have noticed that there are few, if any, longitudinal studies that help us understand how expectations are changing over time and in different sectors of the population. Many citizens now expect data to be shared by new digital services, many other citizens feel threatened by these services, others might have no change in expectations as they do not use them due to lack of skills, access, or desire.

The new stewardship body should establish a model for continually engaging with, understanding and measuring UK citizen's expectations. This will help the body make decisions and help us all understand whether or not we are making progress in a direction that lets humans flourish.

Long-term (+ 2 years)

1. Revisit the principles

The principles recommended in the review are a useful starting point but will need to be revisited. The process for revisiting them will need wide engagement with public sector bodies, businesses and with people in their multiple roles as citizens, consumers and creators. This creates the opportunity for a wider public debate about the future options for how data interacts with the UK's social contract.

2. Recommend new legislation

The UK is going through a period of significant change for data-related legislation with the recent passing of the Investigatory Powers Bill and Digital Economy Bill, and the imminent introduction of the EU General Data Protection Regulation, the UK Data Protection Bill, and the UK's expected exit from the European Union. As these changes settle down and their effectiveness, or not, is understood then it will be time to establish what is needed next.

We suspect that the next wave of legislation will move from an individualistic approach to data rights to combining elements of group access rights, such as those in the 2016 French Digital Republic bill or our ideas on legislation for data infrastructure, with some of the emerging academic thinking that group privacy may be a necessary next step to supplement current data protection legislation.

Depending on the body's earlier research into devolution then the next wave of UK legislation may also include steps to devolve some data powers to more local institutions.

3. Demonstrate global leadership

The data stewardship body should help the UK demonstrate global leadership in data. We should see more individuals and organisations using data to inform their decisions. We should see data being more open and creating more economic and social value. We should see greater trust in data by individuals, due to greater engagement and more ethical and equitable outcomes.

Success will be measured both through seeing data improve the lives of UK citizens but also by seeing other countries following the UK's lead.

Maeve Walsh

Corsham Institute

The Corsham Institute (Ci) is working “to build a fair, inclusive, prosperous and creative society, based on trust and security”. We believe that this can only happen if the citizen is placed at the heart of all decisions relating to the use, application and regulation of technology. As our society becomes ever more connected and the pace of technological change accelerates, this requires constant vigilance, openness and an ongoing dialogue between the technology companies, government, public and not-for-profit sectors, and the public.

There is an opportunity for the new stewardship body to place itself at the centre of this dialogue and establish data governance frameworks that protect the rights and opportunities of individuals, equally, across all parts of society. These frameworks should not just be for citizens, but designed in consultation and collaboration with them.

The establishment of a new body presents huge opportunities to start a transparent, inclusive and comprehensive discussion on some of the challenges posed by the uses of data, both now and in the future. Indeed, an open, public-facing approach will be a prerequisite if the new body is to establish the authority it needs to function in the future: to address complex ethical and legal issues while earning and retaining the trust of individuals.

Putting the citizen at the heart of its objectives

The impacts of an increasingly data-driven society are already felt in individuals’ lives in many ways. For example:

- the information collected about them from their online activities and connected devices, and used to target (or manipulate) the content they receive through search engines, news and social media outlets, and other platforms;
- the unseen data analytics, and unconscious human biases, that shape algorithmic decision-making: whether to sell products online, decide on an individual’s creditworthiness or employability, inform a medical diagnosis or, in future, take control of their driverless car; and
- the implications that the collection, sharing, trading and analysis of data has for individual’s security, privacy and anonymity; and the related data protection and cyber security risks arising from everyday online transactions and the connection of IoT-enabled devices.

All of these scenarios raise complex legal and ethical questions that will take time to resolve. But as policymakers, regulators and others work through the complexities, there is a risk of widening social and economic divisions through unequal access, application and understanding of data-driven services and products.

Establishing the new body

As the body establishes itself, it needs to engage the public, including difficult-to-reach sections of society, in its aims and to demonstrate clearly the relevance of its work to their daily lives and to their futures. The manner in which it sets itself up will determine its long-term success. From the outset, it must be outward-facing and committed to openness. Given that issues around data are complex and often difficult for individuals to grasp, and that faceless organisations don’t often readily engender public trust, the stewardship body will need to work extra hard to make itself relevant and transparent.

It needs also to engage representatives from all sectors in its design, principles and mode of operation so that the evidence and advice it can call on is diverse, and that its medium-to-long term work programme is informed, flexible and responsive.

Short, medium and long-term priorities

Public engagement and building trust should be at the heart of its approach. For example, in the short term it should:

- Design and develop a high-profile, accessible campaign to set out the issues in terms relevant to individuals' daily lives and activities; and run national and online engagement and outreach events to test public attitudes to data use and governance and shape future principles
- Design and deliver a public consultation, informed by the engagement activities and early input of sector specialists, to: agree the Body's principles, functions and approach, including referral, accountability and recourse mechanisms; seek buy-in to the priorities for its medium-term programme; define the inter-relationships with other bodies in the governance landscape; and further increase and deepen public understanding of the issues.

In the medium-long term, it will need to ensure ongoing, clear communication of the stewardship body's principles, activities and work programme and maintain open channels for public engagement, referral and recourse.

Cross-sector collaboration will also need to be embedded from the outset into its structures and decision-making. This should include: recruiting specialists and practitioners from across all sectors who understand the specific management, applications and future developments relating to data use in their areas. An advisory board should be established, comprising legal and ethical experts and representatives of existing regulatory bodies, academics, civil society and consumer groups.

In the medium-long term, the body will need to develop and disseminate clear principles and supporting technical, regulatory and organisational responsibilities to address its agreed work programme priorities; and a robust system-wide infrastructure to ensure they are applied and adopted across all sectors.

These foundations (public trust and transparency, and cross-sectoral representation) are pre-requisites if the body is to be successful in defining and addressing the complex ethical, technical and regulatory issues in a truly cross-cutting way. In the short term, it might lead a review of the regulatory and enforcement landscape to identify gaps and achieve consensus on any extension of the role of the stewardship body to fill them. In the medium-long term, it should undertake an annual review of the priorities for cross-sectoral investigation and resolution, supported by public engagement and consultation.

Opportunities and challenges

Recent work undertaken by the Corsham Institute, as part of its 2017 Thought Leadership programme, has highlighted some of the current challenges in this landscape¹⁶; in particular, how the speed of technological, data-driven change is such that it is increasingly difficult for the policy and regulatory landscape to keep pace while, at the same time, there is a growing lack of public trust in the organisations that handle our data.

This has a number of far-reaching social and economic consequences, not least in relation to fair and equal access to new services, the success of disruptive business models and inclusive economic growth, and the opportunities for improved community and civic engagement.

The new stewardship body has the potential to shape and lead the debate on the kind of data-driven society and economy we want to create both now and in the future, to protect and represent the rights of individuals and the wider good of society, and to challenge all players across all sectors to keep the impact on the citizen at the heart of all their data governance decisions.

About Corsham Institute (Ci)

Ci is a not-for-profit organisation that is working for a fair, inclusive, prosperous and creative society based on trust and security.

16. www.corshaminstitute.org/research

Marko Balabanovic, Michele Nati, Anat Elhalal, Phil Young and Paul Galwas

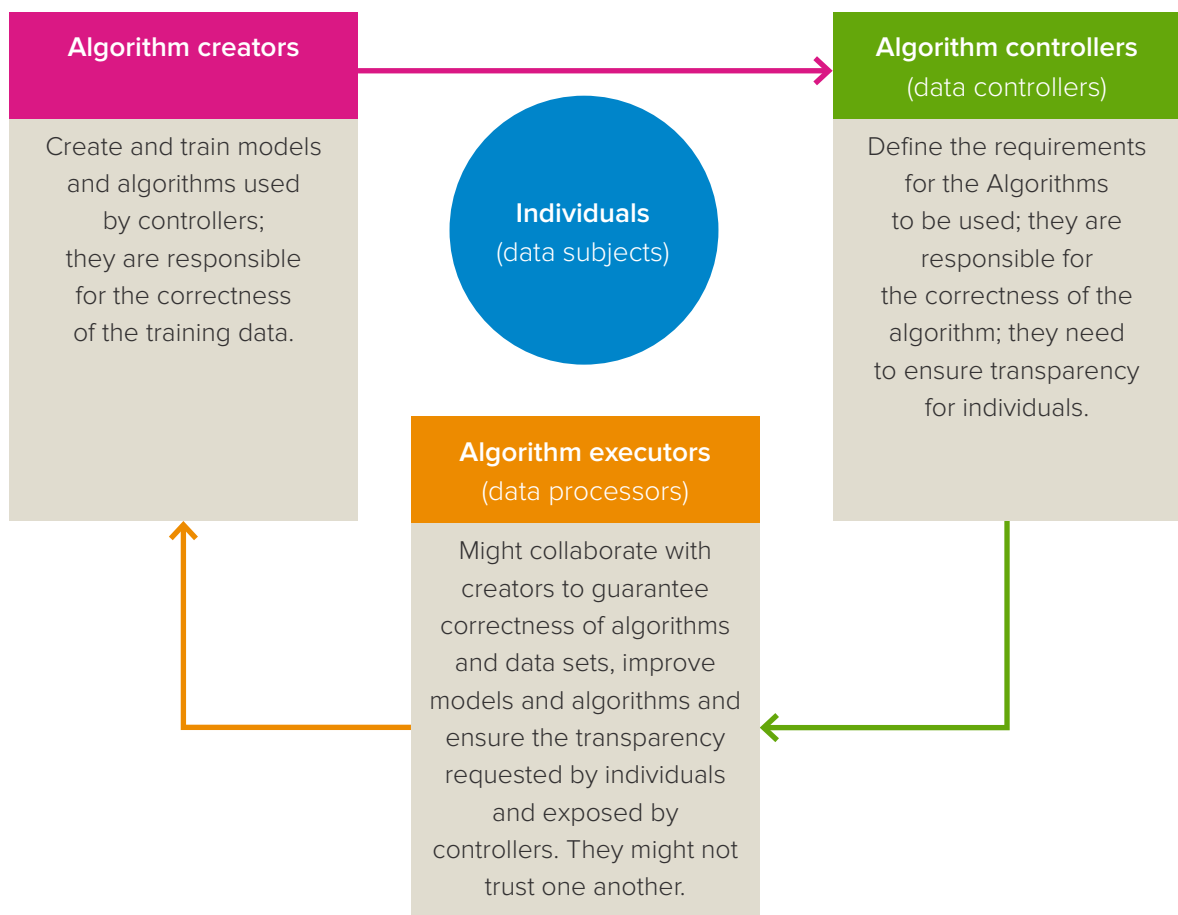
Digital Catapult

Considerations for a data governance stewardship body in the era of Artificial Intelligence

The growing adoption of Machine Learning (ML) and Artificial Intelligence (AI) is increasing the complexity of decisions that machines can take and also the range of contexts in which such decisions can be applied. Moreover, the upcoming General Data Protection Regulation (GDPR), with rights concerning ‘automated decision making and profiling’ is aiming to re-introduce individuals into the loop of data collection and decision making. This poses new challenges for an overarching governance stewardship body.

The situation is further complicated by the emergence of different actors. They span from (1) individuals, contributing data and/or being the subjects of decisions; (2) model and algorithm creators, conducting algorithm training, modelling and development; (3) algorithm controllers, analogous to data controllers in determining the purpose and manner in which the algorithms will be used and (4) algorithm executors applying the algorithms or models to specific data and usage contexts analogous to data processors. While (1), (3) and (4) were already present in the traditional data processing model, already covered by the Data Protection Act as data subject, data controller and data processor, (2) represents an emerging new category.

Figure 1



This complexity of relations introduces new challenges for governance to ensure transparency, trust and assign liability among stakeholders, while guaranteeing compliance.

To understand the range of responsibilities of proper data governance in the case of AI-driven processes, the diagram below shows how the information feeding into AI/ML algorithms is captured and used during the lifecycle of an associated AI-driven service, retained and afterwards disposed of.

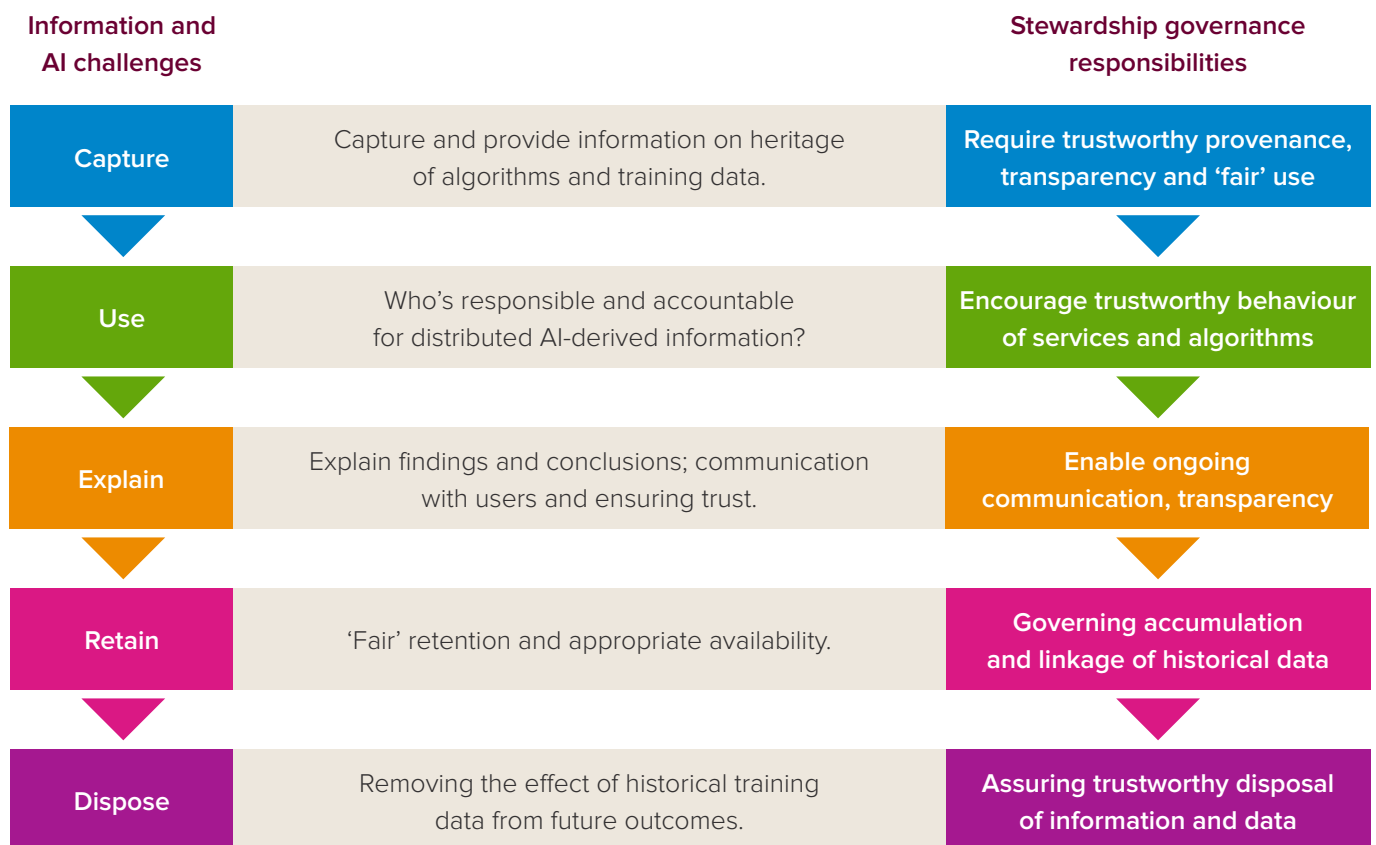
We highlight the new challenges posed by the usage of personal information at the different stages of AI/ML solutions, by comparing them to previous simpler data workflows.

We also identify the role and recommend responsibilities for a new governance stewardship body to consider.

It is worth noting that we consider it vital that a stewardship body explicitly addresses the issues of jurisdiction and scope, national versus international. For example, ‘protect individual and collective rights and interests’ varies from culture to culture, from nation to nation, yet the internet generally does not conveniently partition ‘data’. In particular, when considering applications of AI and ML, we expect complex constellations of cloud providers, multinational organisations and deep learning technology stacks, where jurisdiction may be unclear.

We suggest the role of a stewardship body should be to actively investigate the edges of this space as it develops, according to the dimensions below, and work with related organisations to form policy, gather evidence and create new schemes, to champion experiments and trials, and to disseminate best practice as it develops.

Figure 2



Capture

This phase refers to capture of data for AI/ML algorithms for training as well as inference.

Challenges: What training data has been used and what input information is used for training models? Who owns the training data and who owns the derived models? Is the provenance of the data captured and recorded properly? In particular, under GDPR and other relevant regulations, was the individual advised that their collected data was to be processed by an automated decision algorithm? Has the option to opt out been offered?

Role of stewardship: Issues of fairness, accountability and transparency relating to AI and machine learning systems form an ongoing area of research, and will often come into play when considering the diversity and provenance of training data and processes for algorithms. Depending on the application area, there may be existing governance, regulations and laws that can be applied. However, given the rapid advance of machine learning techniques into hitherto unanticipated applications, there will increasingly be grey areas where existing regulation is unclear, or where ethical dilemmas arise. In the context of the capture stage this must include the diversity and provenance of training sets as well as the transparency of how personal data is used. It should also suggest best practices in providing accurate and trustworthy audit trails, and help ensure that the value created by the use and re-use of derived algorithms is always shared fairly.

Use

This phase refers to how data is used to derive knowledge and take decisions. It considers if (in the case of personal data) a user's will is considered and compliance guaranteed (eg, adequate transparency on automated decision making and profiling); it should identify who is responsible for training the algorithms (ie, Algorithm Creators) and who is actually using the algorithms (ie, Algorithm Executors).

Challenges: Have the user's rights been considered? Are there guarantees of integrity? Are the effects of wrong decisions carefully considered? Is any associated liability properly identified? Are there safeguards against biased decision-making? How is performance and compliance assessed in an ongoing way, particularly of ML systems that adapt over time or react to new situations? Is the risk of reidentification considered?

Role of stewardship: It should monitor that the algorithms' effects are carefully considered by the algorithms' controllers and executors, and that the safety of individuals is a priority, as well as their anonymity in case of anonymised data sets. A stewardship body should create frameworks for auditing and accountability to cover real-world cases for the use of AI and ML algorithms, by ensuring that liability is identified in case of complex relationships such as those underpinning modern AI-based solutions.

Explain

This phase refers to the capability to explain the different findings and conclusions that algorithms produce; their communication to individuals and how to engender trust.

Challenges: How can users be sure that algorithms are fair? How is it communicated that they have been validated and tested? Is it clear AI or ML systems are being used? Are algorithms static or continuing to learn with exposure to new input? What kind of explanations are provided to individuals wishing to know more?

Role of stewardship: A stewardship body should ensure governance of best practices around communicating decisions and means of accessing explanations, ensuring that diversity and accessibility are always considered. In doing so, a governance body should ensure that balance and trust are maintained between all four groups considered earlier. A stewardship body should also facilitate the labelling of algorithms, to provide a widely accessible human-readable communication tool. This could rate algorithm transparency and decision making, while addressing versioning issues.

Retain

This phase refers to making older version of algorithms (including training data sets) available for future inspection and enabling individuals to explore their behavior or to challenge it.

Challenges: Are algorithms' outputs reproducible? For how long should training data, input data and model versions be retained? How should access to this information be given to individuals?

Role of stewardship: A body should address best practices on governing the accumulation and linkage of information managing decision making records, on-going access. A governance body should create a framework for data and algorithm retention and be able to identify and track responsibilities among algorithm creators, controllers and executors.

Dispose

This phase refers to the procedures related to algorithm and model disposal, when new versions are available, or previously available training data sets are not available anymore. It covers different scenarios and ramifications created by either creators or controllers when models or algorithms are disposed, as well as by individuals asking for their data to be removed, when the legal basis for the right to be forgotten applies (such as under the GDPR or similar regulations).

Challenges: How is ownership of training data and the derived model handled in case of algorithm and model disposal? Can deletion of training data be requested, or specific records used in the creation of models? How does this affect the system and other users?

Role of stewardship: A new body should help resolve ownership issues on training data, in particular in the case of a personal information, and how the individual's right to be forgotten might affect the algorithm's behaviour. A stewardship body should set best practices for frameworks for auditability to be provided by creators. It should suggest adequate mitigation strategies for contexts where algorithms are no longer applicable or are removed from service. A governance body should suggest guidelines for communication among all involved parties, creators, controllers and executors when disposal occurs.

Conclusion

In this short paper we explored the challenges and possible roles for a stewardship body by identifying the following phases of data and model development in the AI/ML service lifecycle: capture, use, explain, retain and dispose. It is crucial that the stewardship body continually publishes its methodologies, standards and industry reviews to increase public understanding of how data is used in the different phases and enhance trust in automated products and decision making. We anticipate that the short term focus and priority of a stewardship body should be in the capture and use phases.



**BRITISH
ACADEMY**

for the humanities and social sciences

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Society's strategic priorities emphasise its commitment to the highest quality science, to curiosity-driven research, and to the development and use of science for the benefit of society.

These priorities are:

- Promoting excellence in science
- Supporting international collaboration
- Demonstrating the importance of science to everyone

For further information

The Royal Society
6 – 9 Carlton House Terrace
London SW1Y 5AG

T +44 20 7451 2500

W royalsociety.org

 @royalsociety

 @theroyalsociety

Registered Charity No 207043

Issued: December 2017 DES5315

The British Academy is the UK's national body for the humanities and social sciences – the study of peoples, cultures and societies, past, present and future. We have three principle roles: as an independent Fellowship of world-leading scholars and researchers; a Funding Body that supports the best ideas, nationally and internationally; and a Forum for debate and engagement – a voice that champions the humanities and social sciences.

For further information

The British Academy
10 – 11 Carlton House Terrace
London SW1Y 5AH

T +44 20 7969 5200

W britishacademy.ac.uk

 @britac_news

 @TheBritishAcademy

 Britacfilm

 BritishAcademy

Registered Charity No 233176