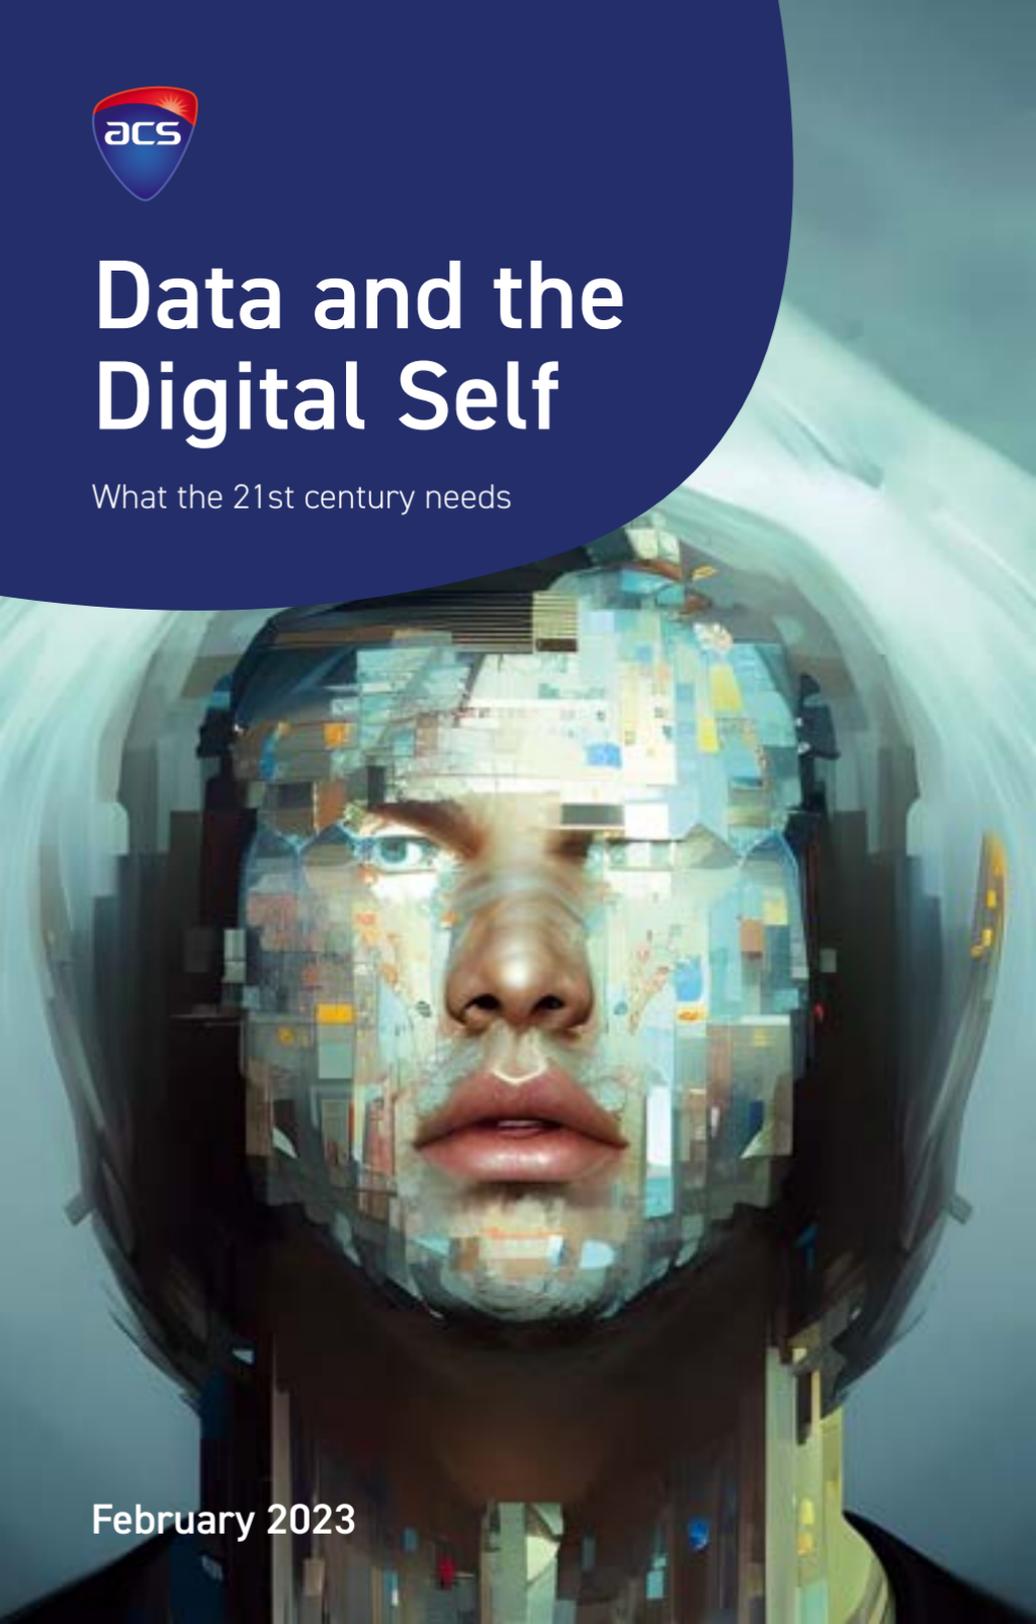




Data and the Digital Self

What the 21st century needs

February 2023





Contents

Chapter 1	2
Introduction – Why we should care	
Chapter 2	4
Our future home – Our future Australia	
Chapter 3	18
Data problems and legal solutions – Some thoughts beyond privacy	
Chapter 4	48
Trust building for data sharing – Understanding trust as a social relationship	
Chapter 5	76
Data privacy, fairness and privacy harms in an algorithm- and AI-enabled world	
Chapter 6	111
Final Words	

Chapter 1

Introduction – Why we should care

It seems inevitable that our future is digital, ubiquitously connected and critically dependent on technology.

For decades, most of the world's economies have been moving towards being service-dominated, and these service economies are increasingly digital, online and driven by data. In all sectors, services are increasingly created, delivered and consumed via digital means, propelled by the increasing adoption of online and mobile technologies.

Recent years have witnessed dramatic changes in the way we consume music, watch movies or arrange dinner. While they may be less obvious, there are dramatic changes underway in all industry sectors, and in government, pushed by changing consumer expectation, reduced barriers to new entrants and an increasingly borderless world of information flow.

Overlaid with our expectation of an ever improving quality of life, major drivers for this change to a digital future world are coming from the intersection of a growing, ageing and urbanising population; a globally changing climate; and a response to an increasing number of global challenges and pandemics. In most economies, these drivers of change are also giving rise to sectoral challenges, such as rising healthcare costs, increasing household energy prices and pressure on government to do more with less.

In a world with finite resources, these challenges must be met with ever greater productivity increases driven by technological enhancements. Continued technological advancement and rapid adoption are central to our ongoing response to this challenge of 'sustainable intensification'.

As technology and digital solutions increasingly play a role in moving the economy and society forward, they become pervasively embedded into business operations, across key service offerings and into our personal lives. New developments spur more innovative business models, products and services that are crucial in responding to our current challenges but also lead to accelerating use of and generation of data and digital services.

By 2030, it is likely to become a self-reinforcing process, accelerated by increased use of artificial intelligence (AI). AI makes sense of the rising tide of data to locally optimise and increasingly personalise service delivery.

By 2030, our dependence on technology will make cyber security crucial to navigating the associated risks and opportunities ahead. This, combined with the growing complexity and sophistication of cyber security threats, makes us more vulnerable at a national, organisational and individual level.

As new systems develop, privacy and consent will need to be a central pillar of the process of collecting, sharing and using these datasets. Methods and frameworks will need to be developed for providing and handling consent, for sharing and using data, and for providing security in highly complex networks.

In this collection of perspectives on data and our future, we explore the challenges of emergence of global trends and technologies, examining how these trends will shape our concepts of future Australia and impact our concepts of privacy, consent and 'appropriate use'.

Chapter 2

Our future home – Our future Australia



By Dr Ian Oppermann FACS
ACS Immediate Past President

Ian currently holds the role of NSW Chief Data Scientist and is an Industry Professor at UTS. Ian has 30 years' experience in the ICT sector and has led organisations with more than 300 people, delivering products and outcomes that have impacted hundreds of millions of people globally. He has held senior management roles in Europe and Australia as Director for Radio Access Performance at Nokia, Global Head of Sales Partnering (network software) at Nokia Siemens Networks, and then Divisional Chief and Flagship Director at CSIRO. Ian is considered a thought leader on digital economies and is a regular speaker on big data, broadband-enabled services and the impact of technology on society. He has contributed to six books and co-authored more than 130 papers that have been cited more than 4,000 times. Ian has an MBA from the University of London and a Doctor of Philosophy in Mobile Telecommunications from the University of Sydney.

Blazing summertime

A little, azure bird shines because of the brook

The Australian Bureau of Statistics told us that there were 309,996 registered births in 2021, an increase of 5.3% from 2020. The significant majority of these 309,996 children are likely to live to well beyond 100 years of age, taking them well into the 22nd century. During the course of their lives of 100+ years, these children will see change unprecedented in our lifetimes. These children will still be pre-teen in 2030, the timeframe that many planners use as an accessible planning horizon. They will be around 30 years old when they hit 2050 and face very significant societal, technological and environmental changes.

As context, a child born today is born into a wider world of an estimated 8 billion other inhabitants.¹ By 2030,² this number will grow to an estimated 8.5 billion, and by the time our child enters 2050, there will be an estimated 9.7 billion people on the planet. During this time, there will be no more land created, no more water produced and no more natural resources beyond what we have already in our closed system. We know the world's population is living longer,^{3,4} and so getting older on average, and moving to increasingly densely populated cities.⁵ These elements are creating a long-term productivity challenge as the ratios of people working to those retired start to change significantly. We also know the climate is changing globally, which impacts where and how we grow our food. These are the major elements of the sustainable intensification challenge.

Some of these global trends will be directly relevant to us here in Australia; some will be less intense because of our unique national circumstances. Nonetheless, like the rest of the world, we need to think carefully about how we use the resources we have, and plan for the changes that are inevitably coming our way.

We begin by focusing on a near horizon, the trends behind change to 2030, and a framework to give us the tools to discuss how that Australia could work, in terms of 'smart' services and places.

1 https://www.un.org/development/desa/pd/sites/www.un.org.development.desa.pd/files/wpp2022_summary_of_results.pdf

2 <https://www.un.org/development/desa/en/news/population/world-population-prospects-2019.html>

3 <https://data.oecd.org/healthstat/life-expectancy-at-birth.htm>

4 <https://www.statista.com/statistics/673420/projected-global-life-expectancy/>

5 <https://ourworldindata.org/urbanization#:~:text=Urbanization%20is%20a%20trend%20unique,areas%20as%20they%20become%20richer>

Some context for the future of Australia

For our child born today, there are major drivers of change coming from the intersection of our growing, ageing and urbanising population, from technological advances, from changes in society, from a changing climate and from global shocks. Data and digital services help us understand these changes as well as develop responses to them.

According to the Australian Government's Centre for Population,⁶ Australia's population is expected to be smaller and older than projected prior to the onset of the pandemic. Australia's population is estimated to be around 4% smaller (1.1 million fewer people) by 30 June 2031 than it would have been in the absence of COVID-19. The population will also be older, as a result of reduced net overseas migration and fewer births. Australia's population is still growing and is expected to reach 28 million during 2027–28, two years later than was estimated before COVID-19.

The move to a digital economy, and digital engagement with others, alleviates the 'tyranny of distance' that has historically provided challenges for community development, access to markets and growth of local industry.

The advent of COVID-19 restrictions showed just how effectively this can be done for many industries and job types. It also showed just how dependent we have become on network connectivity. Those regions with limited connectivity were least able to adapt to living and working online. Those industries and job types that were least able to move online remain significant employers and drivers of Australia's economy. The blend of online and real-world interactions and value creation is nonetheless likely to move towards increasing use of online, digital services in all parts of Australia.⁷ Widespread access to reliable network connectivity therefore becomes an even greater long-term consideration for all communities.

Although Australia has a long history of cycles of fire and flood, we have seen the growing scale of the impact on communities. Towns in Queensland and New South Wales that needed to truck in water to keep communities going in 2021⁸ were the same communities struggling to cope with a series of flood events in 2022. Future developments that ignore the impact of a changing climate – whether it be water availability or the growing intensity and frequency of natural disasters – do so at their peril.

To understand the impact of natural disasters on communities and infrastructure, and to deliver effective responses to disaster events, requires increased use of data in many forms, potentially from traffic movement to economic activity. It's also important to align with the national response to climate change. The Australian Government has committed to reducing greenhouse gas emissions by 43% below

6 <https://population.gov.au/publications/statements/2020-population-statement>

7 See, for example, ACS Australia's Digital Pulse 2021. <https://www.acs.org.au/insightsandpublications/reports-publications/digital-pulse-2021.html>

8 See, for example, <https://www.abc.net.au/news/2021-01-10/southern-queensland-still-in-drought-while-north-floods/13039008>

2005 levels by 2030 and net zero by 2050.⁹ This requires cities, towns and places to consider and measure their own contributions to these targets.

One major area of response to these challenges is the ongoing focus on making our cities, towns and places 'smart', which is underpinned by the use of data. Our response to COVID-19 clearly demonstrated the value of access to data for rapidly responding to the pandemic. We needed to understand the impact of travel and other lockdown restrictions on critical services and critical infrastructure – from ensuring nursing homes remained safe to ensuring power stations remained in operation. This required a new capacity to understand a city in near real time, and to have realistic, data-driven models that allowed options to be explored.

Access to unprecedented datasets from mobile communications¹⁰ to credit card transactions – all in aggregate form to protect individual privacy – allowed governments to understand the effectiveness of health order restrictions on movement, and the economic impact of these same restrictions as well as the subsequent economic stimulus. The use of these aggregate, people-centric datasets was an important element in governments' response to COVID but raises the issue of the inherent need to create and maintain trusted frameworks to use these data for agreed (and important purposes).

Despite all of these challenges, and in the face of our changing population profile, we continue to expect an ever-improving quality of life, or at least to not go backwards. These are challenging factors to reconcile. They must be met with broad outcomes-based thinking that clarifies what we are trying to achieve, how we can tell when we are achieving the outcomes, and the means to understand why we are not.

It also requires us to make our services smarter, whereby we can understand in fine detail what is happening in a city or community, identify root causes of problems, and even be able to predict when things will go wrong and plan adaptive scenarios to respond to changing needs. This relies heavily on access to data, and this access to data will be influenced by technology trends around the creation of and use of this data. It also opens up questions on privacy, security and consent.

Technology trends shaping a future Australia

For decades, Australia has been transforming towards a service-dominated economy¹¹, and our service economy is increasingly digital, online and driven by data. Continued technological advancement and rapid adoption are central to national progress and are crucial aspects of our responses to the challenges highlighted in the previous section. Both will be an increasingly important part of our future Australia.

9 <https://www.industry.gov.au/news/australia-submits-new-emissions-target-to-unfccc>

10 See, for example, <https://www.themandarin.com.au/129871-federal-nsw-governments-use-vodafone-data-to-see-if-public-is-following-covid-19-restrictions/>

11 www.abs.gov.au/articles/services-australian-economy

Efforts to create joined-up customer experiences make it easier to engage with service providers and hopefully deliver more effective outcomes. For many years, businesses have tried to create a better, more attractive customer experience, delivering increasingly personalised services. The effort by NSW and other governments to create 'life events' information services¹² is aimed at making (unavoidable) engagement with government more seamless for important events such as getting a job, dealing with the death of a loved one or the birth of a child. In all cases, the exchange of people centric data across different tiers of government helps make the engagement easier, delivering more of the total solution required and leaving less for the citizen (or customer) to follow up themselves.

New services and capabilities promised by access to data and digital technology present enormous benefits for users and citizens – delivering services and information to people when they need it, wherever they happen to be, whenever they need it and in whatever circumstances.

Inescapably, however, any data collected about people directly – their actions, location, environment or any aspect of the context they operate in – has some aspect of what may be regarded as personal information, even if 'de-identified' to remove unique identifiers. If the datasets used for these purposes are linked and analysed to provide sophisticated, personalised services, a great deal of personal data (PD) or personal information (PI) may be contained in the joined data, possibly sufficient to re-identify the individuals represented therein.

Let us look at some technology areas in more detail and consider the consequences for planning smart places today.

The march of the Gs

The introduction of widely available mobile communications in the 1990s in the form of GSM (2G) fundamentally changed our world, from the way we socialise to how we conduct business. Many would recall the liberation of being able to make or receive a call from (almost) anywhere. Since then, technology has marched along, with 3G in the early 2000s, 4G in 2010 and 5G from 2020.

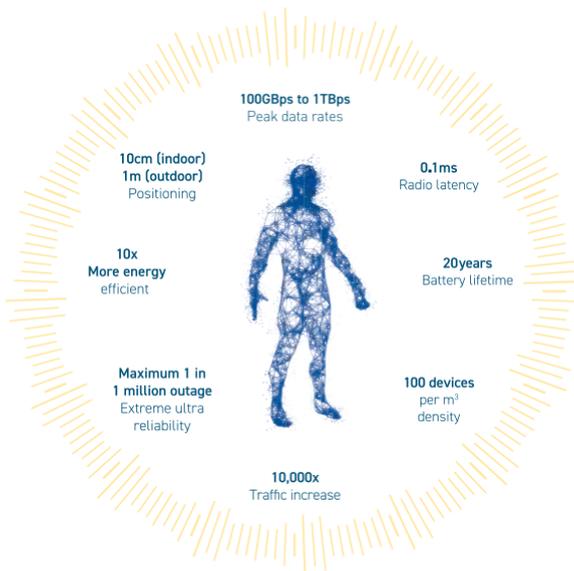
As the technology has improved, the dominant use case of making or receiving calls has increasingly been replaced by access to the internet, data services and connections to devices rather than people. With 5G came the ability to reliably monitor and operate remote devices, to stream multiple channels of high-definition video, or drip-feed a few bits at a time between millions of sensors measuring anything from soil moisture to air quality. Voice calling remains a feature of 5G but is now a 'special case' functionality rather than the main driver. Movement of data is now the major function of mobile networks.

¹² See <https://www.nsw.gov.au/life-events>

The race towards developing 6G by 2030 has begun. It will underpin widely anticipated future services – from immersive entertainment to secure and reliable operation of autonomous devices, and from monitoring personal health devices to securely and reliably managing driverless cars, trains, ships and aircraft.

All these services are critically dependent on three combined elements. The first is instant, virtually unlimited wireless connectivity, which combines with the second, access to a wide range of constantly developing and potentially highly personalised datasets. They move 6G towards a distributed 'prosumer' network where data is constantly generated as well as consumed by participants in the network.

The third important element is widespread use of AI to make sense of the rising tide of data to generate insights, to spot anomalies and to locally optimise systems. AI can also be used to augment human systems through direct human engagement (personalised assistants) or as intelligent autonomous mechanical systems (classical robots). These three factors combined – connectivity, data and autonomous intelligent systems – create a range of new considerations when we contemplate the trade-offs between the uses of data for delivery of highly individualised services and treatment of personal information throughout the network, ensuring consent is obtained and handled in a meaningful way, and developing security frameworks in highly complex systems. These considerations are present for existing networks; however, they are fundamental to the envisaged goals of 6G.



General requirements for 6G

Source: Latva-aho M and Leppänen K (eds) (2019) Key Drivers and Research Challenges for 6G Ubiquitous Wireless Intelligence. <http://jultika.oulu.fi/Record/isbn978-952-62-2354-4>

Data as a driver for service innovation and privacy

The world has seen waves of technology-driven innovation in all sectors from finance and telecoms to government services and smart infrastructure. In these environments, data is one of the most significant transformative factors, creating a means of increasing transparency, supplying a source of innovation, and providing the ability to understand, optimise and personalise information and services. The ability to harness a wide range of large, constantly evolving and highly personalised datasets is a strong source of productivity and supports the creation of new, high-value services. The 'smart' in smart services comes from accessing and using data.

Data use also requires consideration of personal privacy. There is currently no unambiguous nationally accepted test for personal data (PD), personal information (PI) or personally identifiable information (PII) in a dataset. Often the terms are conflated. Most privacy assessments worldwide rely on tests of judgement described in terms such as 'reasonably' or 'likely'. The Commonwealth Privacy Act 1988¹³ defines personal information as:

Information or an opinion about an identified individual, or an individual who is reasonably identifiable:

- a. whether the information or opinion is true or not; and*
- b. whether the information or opinion is recorded in a material form or not.*

If datasets are linked and analysed to provide rich new services, a great deal of PD or PI may be contained in the joined data, possibly sufficient to re-identify the individuals represented therein. Most privacy assessments rely on tests of judgement described in terms such as 'reasonably' or 'likely'.

The challenge is to quantify the amount of PD or PI in a dataset at any point in time and in any given context. This extends to developing threshold tests for when an individual is 'reasonably identifiable', while considering personal attributes, temporal and spatial aspects of data, and rich contextual environments. Some of these challenges are yet to be fully addressed. In the meantime, guidance is available from the Office of the Australian Information Commissioner (OAIC)¹⁴ and each jurisdiction's Privacy Commissioner or Information Commissioner.

The consequences for future services design include:

- ensuring that every person involved in smart services activities has an appropriate level of data literacy, an awareness of privacy legislation relevant to the jurisdiction, and an awareness of data governance fundamentals
- developing data governance frameworks for data use as part of smart services project design
- adopting a 'privacy by design' mentality¹⁵ with project development

13 <https://www.legislation.gov.au/Details/C2014C00076>

14 <https://www.oaic.gov.au/>

- ensuring there are clearly identified roles and responsibilities in data sharing and use scenarios, including identified privacy governance experts with deep capabilities
- reviewing projects throughout their life cycle for privacy and data governance considerations.

Increasingly smart, connected and complex – Creating a cyber security challenge

The complexity of systems stems from demands to improve efficiency and effectiveness, promote productivity and adapt to changing consumer demands. Multidirectional energy flow from prosumers contributing to smart grids supports the use of renewable energy sources in national grids and supports important future capabilities such as meaningful demand-side energy management. It also takes the traditional unidirectional network model into a complex new direction.

This creates significant new technical challenges, the need for data interchange between multiple stakeholders, and the need to assure the system will operate correctly in a wide range of new situations. This is true of any newly 'smart' system, from telecommunications to smart traffic. The active component in these smart systems has greater potential to impact the outcome and become a point of vulnerability in a system.

Data will increasingly become the most valuable asset in digitally enabled systems, especially data that contains sensitive personal and commercial information. For consumers, businesses and government to increasingly gain confidence in digitally enabled systems, and so increasingly rely on such systems, data must be protected from misuse.

The protection of data should not be limited to when it is at rest (such as when it is stored in a disk or memory) or in transit (when it's being transmitted in a network), but also when it is used to model, plan and optimise services. This is increasingly important for smart critical infrastructure from telecommunications to smart grids. As technical functionality increases, the cyber security requirements escalate dramatically.

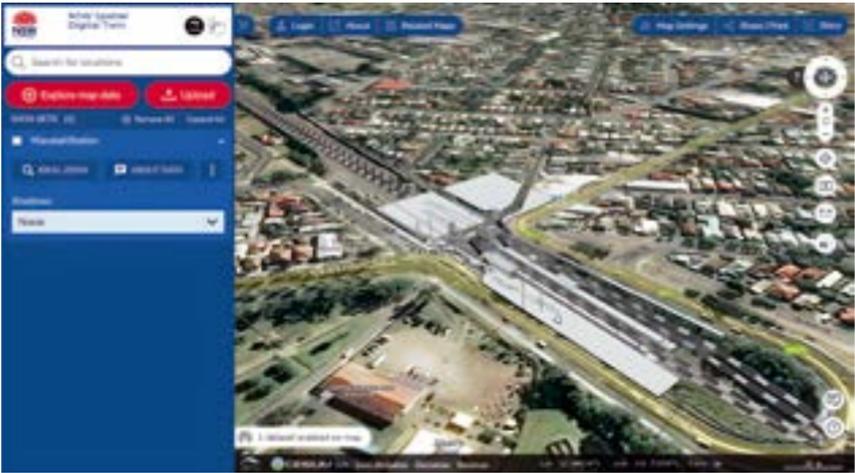
For some time, cryptographic approaches have been used to protect data where the data is encrypted both in motion and at rest, so that they are never revealed to anyone other than data owners themselves. However, searching and processing encrypted data can be extremely inefficient and costly when it requires transfer to a trusted server for decryption and processing.

15 See, for example, OAIC. <https://www.oaic.gov.au/privacy/privacy-for-organisations/privacy-by-design/#:~:text=Privacy%20by%20design%20is%20a,of%20new%20systems%20and%20processes>

Digital twins

A digital twin is a virtual representation of a physical object or process; these are used in fields as varied as manufacturing, health, entertainment, construction and geological modelling. The digital twin connects to the physical object or process through a constant stream of updates (possibly in real time), which maintains the consistency of the physical and real worlds. Many digital twins are proprietary, specifically developed for their environment.

A spatially enabled digital twin combines a digital twin with spatial and positioning information, covering a defined geographic space. Spatial digital twins build on the concept of building information modelling (BIM), which has recently been standardised.¹⁶ Very ambitious spatial digital twin projects are under development in different jurisdictions in Australia (including the NSW Spatial Digital Twin shown) and they allow a number of planning and construction scenarios to be explored before real world deployment.



NSW Spatial Digital Twin

Source: NSW Spatial Services, https://www.spatial.nsw.gov.au/what_we_do/projects/digital_twin

The Gemini Programme of the Centre for Digital Built Britain's National Digital Twin programme (NDTp)¹⁷ released the Digital Twin (DT) Toolkit. This includes a guide outlining the key areas of consideration in the approach to developing digital twins.

¹⁶ See ISO 19650-1:2018 – Organization and digitization of information about buildings and civil engineering works, including building information modelling (BIM) – Information management using building information modelling – Part 1: Concepts and principles at <https://www.iso.org/standard/68078.html>

¹⁷ <https://www.cdbb.cam.ac.uk/what-we-do/national-digital-twin-programme>

Artificial intelligence

AI is already used by almost every Australian in some form or other, from the form of smart personal assistants to wayfinding by navigation systems, or even interacting with chatbots. AI's long and steady climb out of movie storylines into real-world use has accelerated in recent years. AI capability is expected to benefit from the increasing availability of ever richer datasets, and from increasingly powerful computing environments leading to accelerating capability and consequently much more pervasive use.

AI will bring greater personalisation of services, augment decision-making, and be used as a frontline defence in escalating cyber security challenges. AI will also help entertain and educate, help identify anomalies in the digital and physical world, optimise systems, and lead to ever greater use of data.

In the legal world, AI is already being used to analyse large numbers of documents through increasingly sophisticated e-discovery, surfacing important insights from vast numbers of digitised documents. 'Rules as Code' is a concept where systems are described in a formal language that is also human-readable.¹⁸ Compiling a regulation or system using formal language allows internal inconsistencies to be identified and allows interaction between different regulations to be explored.

One of the major considerations for a future world is the extent to which AI and automation are used. Acceptance of automation is framed within a wide set of concerns including unintended consequences of automated decision-making, the need for human judgement in the decision-making process, concerns about loss of jobs and even ethical considerations.

NSW has developed and released its Artificial Intelligence Strategy,¹⁹ AI Ethics Policy and AI Assurance Framework, an implementation guide to assist agencies to bring AI projects to life in a responsible manner. The NSW strategy makes direct reference to the international standards community in the form of ISO/IEC JTC1/SC 42.²⁰

Increasingly personalised

Personalisation of services is something we have all come to expect. When we access social media sites, we expect to be offered news and information that is relevant to our interests. Often, when we sign up for a new service, there is a short series of questions we answer to help shape that personalisation. We engage willingly to kickstart the personalisation process.

Very often, however, the systems we use continue to learn about us and our preferences as we continue to use them. Consumers may even express frustration if services offered are not tailored or require the same information to be provided on

18 See, for example, <https://oecd-opsi.org/projects/rulesascode/>

19 <https://www.digital.nsw.gov.au/policy/artificial-intelligence-ai/ai-strategy>

20 <https://www.iso.org/committee/6794475.html>

multiple occasions. The challenge for service users is understanding the amount of personal information that has been shared and the limited ability to control future use of that shared information.

Personalisation also has advantages for those delivering the services. Understanding preferences and the trade-offs a user is willing to make allows for local optimisation of resources. For instance, a customer's willingness to delay use of electricity in exchange for lower 'spot prices' for that power provides the potential to more readily match peak energy availability with peak demand. This can have substantial impact on infrastructure costs.

Similarly, understanding a customer's willingness to briefly accept a lower grade of mobile phone connectivity in exchange for lower communications costs can also help manage scarce resources at a local level. The challenge for the providers is how to gain this understanding of preference and willingness to engage in these trade-offs.

The very nature of locally optimised and delivery of highly personalised services create these challenges. Taking advantage of these services requires considering new methods for providing and handling consent, new frameworks for sharing and using data, and new issues for security in highly complex networks.

One of the principal technical challenges will be to develop a measure of the level of personal data (PD) or personal information (PI) in datasets used for the delivery or optimisation of services, as well as to determine thresholds for when this personal information measure exceeds the 'reasonable likelihood' of identifying an individual. The measure of PI must go beyond simply considering personal attributes captured in data and must consider preferences revealed through the use of services, and temporal and spatial aspects of data, as well as context for the use of services.

A digital challenge – Consent

Just as the proliferation of connected devices creates new considerations for privacy, similar issues arise when considering 'consent', referring to the explicit, informed and freely given consent for data about an individual to be created, transmitted, stored, used and re-used. For example, under Article 6 of the European General Data Protection Regulation,²¹ businesses must identify which one of the six possible legal bases allows their data processing. Consent is one of these six and is one of the easiest to satisfy as it allows businesses to undertake a wide range of uses of the data, provided what is to be done is clearly explained and explicit permission is obtained from the data subject. In Australia, guidance is given by the OAIC as:

You give express consent if you give it openly and obviously, either verbally or in writing. For example, when you sign your name (by hand, or by an electronic or voice signature). An organisation or agency must get your express consent before handling your sensitive information.

²¹ <https://gdpr-info.eu/art-6-gdpr/>

If a human subject of data provides explicit consent, then a human operator or owner of a device can use that data for the stated purposes. The challenges arise when data about a person is collected inadvertently, when the human subject is unaware of data collection about them, or when the human operator or owner of a device is unaware of data being collected about an individual. It may also create problems from a consent perspective when data about an individual is initially used for service delivery then re-used for other purposes such as local optimisation.

Consider the not-too-futuristic example of a swarm of drones providing local mobile hotspot coverage during peak hour commute times. One drone may move to follow the driverless vehicle of a high-data-use individual who is engaged with in a multiparty videoconference during the commute. To provide effective hotspot coverage, the drone must track the driverless vehicle relatively closely (or make a staged handover to other drones) giving the drone information on origin, destination, route taken, time of day and volume of data traffic used by the data subject. If the drone is also used for other non-telecommunications purposes such as vehicular traffic control and environmental monitoring, it may regularly report the data subject's vehicle location as part of a city-wide traffic profile, and even monitor ambient temperature, possibly including the temperature of the passengers in the autonomous vehicle.

The level of PD or PI which has been accumulated will depend on many factors including duration of the coverage by a drone (or network of drones), the spatial and temporal resolution of vehicle location reporting, and the resolution of temperature monitoring. It also depends on the form in which the data has been collected, transmitted, stored, analysed and ultimately used. In all cases, it is unlikely that the data subject would expect to need to consent to their body temperature being recorded as part of the agreement to access mobile communications services.

If, however, the ability to (possibly inadvertently) capture body temperature was considered as part of the informed consent of the user to enter a discounted mobile network plan and if a number of protective measures were taken by the service provider – such as the non-telecommunications data being wrapped around with protective governance, the raw data never being seen by a human person, raw data only ever being processed by machine, insights assured to only be released in aggregate, and any harms that arose to be met with swift compensatory actions by the provider – then this trade-off of use of personal data may be given with reasonable consent.

Data sharing frameworks – Making ‘smart’ in a trusted way

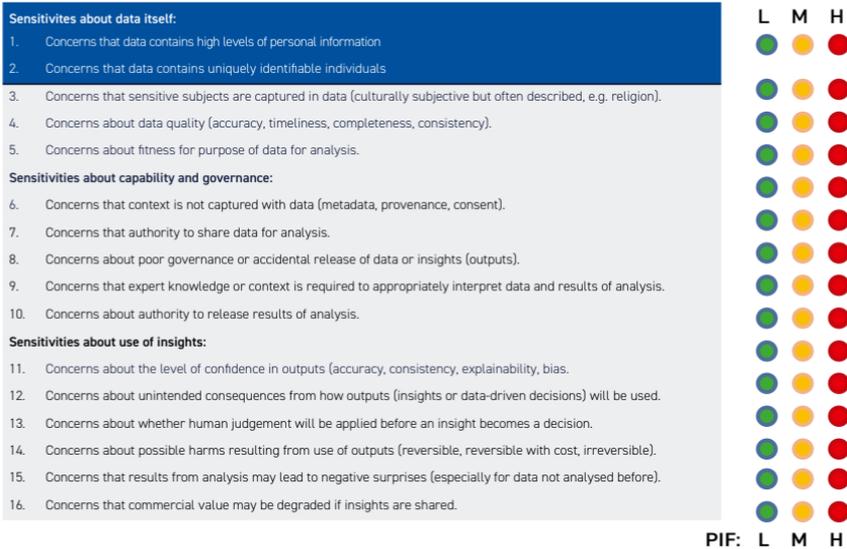
The ‘smart’ in services comes from the ability to access and analyse data to improve situational awareness, understand or even predict root-cause challenges, deliver high-value new services and explore different possible scenarios.

Frameworks are required to collect, share and analyse data for smart city services, and for the application of sophisticated analytics to understand the data. It creates a range of new considerations when we contemplate the trade-offs between optimisation of systems or personalisation of services, and personal information, consent and security.

Considerations for data use

Often objections to data sharing and data use are framed in terms of PD, PI or PII and the requirement to adhere to local privacy legislation. The actual concerns of data holders, however, often relate to sensitivities associated with the data itself, or consequences of use of the data.

Being able to identify the inherent sensitivity of data and different levels of personal information at each stage in the data life cycle provides an opportunity to develop different data governance and data handling frameworks. The figure below highlights a number of considerations (or concerns) across the data life cycle when considering data sharing and use.



A sample of concerns related to data sharing and use

It is important to consider the entire life cycle model when considering data sharing or use. A generic model that can be applied to data sharing and use will include stages where authorising frameworks are assessed and confirmed, and metadata on conditions of capture or impact of transmission and use are recorded, along with provenance and chain of consent information.

Concluding remarks

We have explored the journey to 2030 using the metaphor of a child born today, looking at aspects of what that world will look like and some of the outcomes we would like to influence.

This child's future is digital, ubiquitously connected and critically dependent on technology.

As technology and digital solutions continue to play a key role in driving the economy and society forward, they become increasingly embedded into business operations, across key service offerings and into our personal lives.

By 2030, this will have become a self-reinforcing process, accelerated by increased use of AI to handle rapidly growing datasets to continue to locally optimise and increasingly personalise services delivery.

By the end of this decade, our heavy dependence on technology will make cyber security critical. Increasingly sophisticated cyber security threats will increase the threat to every level of society.

Future services have the potential to deliver enormous benefits; however, their very nature highlights challenges when contemplated within existing regulatory frameworks. As new systems develop, privacy and consent will need to be a central pillar of the process of collecting, sharing and using these datasets. Methods and frameworks will need to be developed for providing and handling consent, for sharing and using data, and for providing security in highly complex networks.

As our child born today grows into their thirties, they will reach the world of 2050. We hope this world will still be shaped by the work we are doing today.

Chapter 3

Data problems and legal solutions – Some thoughts beyond privacy



By Lyria Bennett Moses

Lyria is Director of the Allens Hub for Technology, Law and Innovation and a Professor and Associate Dean (Research) in the Faculty of Law and Justice at UNSW Sydney. She is also co-lead of the Law and Policy theme in the Cyber Security Cooperative Research Centre and Deputy Director, Law and Justice in the UNSW Institute for Cyber Security. Lyria's research explores issues around the relationship between technology and law, including the types of legal issues that arise as technology changes, how these issues are addressed in Australia and other jurisdictions, and the problems of treating "technology" as an object of regulation. She is on the NSW Information and Privacy Advisory Committee, the Executive Committee of the Australian Chapter of the IEEE's Society for the Social Implications of Technology, and is a Fellow of the Australian Academy of Law.



and Kimberlee Weatherall

Kimberlee is a Professor of Law at the University of Sydney, focusing on the regulation of technology and intellectual property law, and a Chief Investigator with the ARC Centre of Excellence for Automated Decision-Making and Society. She has extensive experience in conducting and leading interdisciplinary research, integrating philosophy, data science, law, and the social sciences. Her current work focuses on data-sharing and data governance, with a view to ensuring inclusive, responsible, ethical use of artificial intelligence and other systems of automation. Current projects include work on socially responsible insurance in the age of artificial intelligence, and the use of automation in government.

Stormy summertime

A beautiful bird flies, sings

enjoying the wood

Introduction

There is general dissatisfaction with the state of privacy law in Australia. Various possible reforms have been touted, among them the adoption of a data protection law along the lines of the European General Data Protection Regulation (EU GDPR). This would involve inter alia restricting secondary uses of data, enhancing transparency and consent requirements, and creating a stronger enforcement and penalty regime.

But the GDPR has also been described as the 'best data protection law for the 20th century'. Many have argued that we need to think beyond improving mechanisms for notice and consent that rely on individuals to manage how their data is used.

We argue that Australia right now has a unique opportunity to rethink what kinds of data and privacy protection we want, and to move beyond current approaches. In particular, we need to abandon the assumptions built into current legal frameworks in light of the changes wrought by new uses of data opened up by developments in machine learning and other AI techniques. Rather than provide a deep analysis of any particular proposal for reform, our chapter outlines some ways in which we might think outside the box.

Our starting point is the question 'What data problems are we worried about?'

We ask what modes of collection and which uses of data cause harms that law ought to address. In the usual policy discussions, these are treated as 'privacy' issues (Leonard 2020) but our concern is less about the specifics of Australian privacy laws and more about identifying and describing our current data problems, or data dilemmas.

Later, we ask whether these dilemmas are addressed in current legal frameworks – chiefly, Australia's data protection and anti-discrimination laws, as well as the AI ethics movement. Having pointed out gaps in the existing framework – gaps we think are too large to be addressed by merely updating that framework – we turn to alternatives *beyond* privacy.

In the final part of this chapter, we describe some alternative, well-recognised starting points for protecting individuals from harms associated with modern data practices: rule of law principles, and human dignity and autonomy. We explore ways forward – reformed legislation, rational and inclusive discourse, and systems built to align with human-centred values.

Data problems

To develop law for 21st century data practices, we need first to describe what our data problems are. This is far from straightforward. In everyday policy debates, data practices that are perceived to be harmful or problematic are often resisted through a demand for privacy. Without saying anything about the importance of 'traditional' privacy concerns, we start with some very modern problems. All of the following phenomena relate to the collection or use of data. All at least potentially raise privacy concerns, because they all concern the use of data about people in ways that affect those people. None are simple.

Consider *personalisation* – of services, products, pricing, and information. The large online platforms collect vast troves of information about individuals, their activities and interests, creating detailed data profiles or 'data doubles' (Cohen 2019) that can be analysed to direct and tailor products, services, pricing, and even information flows. This use of data can have many positives (Centre for Data Ethics and Innovation 2020).

For individuals, it can bring efficiency in the form of more relevant information, search results and advertising, and better tailored products and services; for business and others, potentially more efficient targeting of services, products and communications.

But personalisation is not without harm. Benign attempts to better tailor marketing messages can rapidly shade into the creepy, such as the commonplace experience of an advertisement popping up that relates to a conversation you've just had. Or the harmful, such as deliberately manipulative techniques to confuse consumers, or determine when people are vulnerable, perhaps depressed, and exploit that weakness (Manwaring 2018). Or the troubling – like 'personalised pricing' meaning people pay very different prices for exactly the same product without the transparency and openness of bargaining in a traditional marketplace. Some price discrimination is commonplace, but it can be concerning and disempowering for the consumer if it lacks visibility or a rationale, especially when higher prices are imposed on historically disadvantaged groups.

At the extreme end, online targeting can even descend into discrimination against people in protected categories. Examples include women being shown different job advertisements, women being selected out in automated recruitment filtering, people of colour not being shown some advertisements for tenancies, and African Americans being more likely to receive a 'false positive' high reoffending risk score that affects how they are sentenced and imprisoned for offences.

These are not imagined scenarios: Amazon shut down an automated hiring system that was shown to discriminate against women; litigation brought by the US Department of Housing and Urban Development in 2019 alleged that Facebook violated US law by allowing advertisers to limit housing ads based on race, gender and other prohibited characteristics; the COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) system has been used to analyse the risk of reoffending in the US, affecting bail, sentencing and probation decisions despite a racially skewed false positive rate.

Even short of outright discrimination, a purpose of personalisation is to treat people differently on the basis of their characteristics, raising challenging questions about which characteristics can justifiably be used for different kinds of personalisation. Should you pay a different insurance premium based on whether your parents separated during your childhood or not, if it can be shown quantitatively to impact risk, although we may not understand why? Are we comfortable with firms distinguishing between people and treating some more favourably, based on factors beyond their control?

Personalisation extends beyond commercial advertising. Increasingly targeted news can mean that we see what is relevant to us, but it can also make us less well informed, as our news is being tailored to confirm our world view because that is what sells. Personalisation of political messaging could segregate the polity, sending people different messages and reducing our overall ability to have a rational debate or to hold politicians accountable for their messaging.

Private actors aren't the only ones gathering more data and using it for analysis, prediction and automation of decision-making. Different levels and institutions of government are increasingly building the capacity to observe citizens, residents and visitors; to gather information about them; to link datasets, and process, analyse, and use that information for general policy development and as it interacts with, and makes decisions about, individuals (Dencik et al. 2019).

In 2021, an intergovernmental agreement on data sharing between Commonwealth and state and territory governments concluded. In 2022, the Commonwealth government enacted the *Data Availability and Transparency Act 2022* to streamline intra-government data sharing as well as sharing with research institutions.

Draft legislation in Australia – the *Identity-matching Services Bill 2019* and the *Australian Passports Amendment (Identity-matching Services) Bill* – would deploy facial recognition technology to match separate identity files including across federal, state and territory jurisdictions and to a more limited extent with the private sector. Through this scheme (currently undergoing a redraft following a recommendation of the Parliamentary Joint Committee on Intelligence and Security), there is the potential for significant amounts of data to be collected and retained, potentially including live feed CCTV and images from social media. In theory this could have significant benefits, including improved convenience in accessing services, reducing identity theft and facilitating law enforcement investigations.

At a general level, more information can lead to better policy analysis and more efficient targeting of limited public resources. Data sharing across government and between levels of government can make public services more convenient, as with the much-vaunted goal 'tell us once': where citizens only need to provide information one time, rather than to every separate department for each individual service for each reporting period

But when collection of data is by the state, we may have reasons to be more concerned about the impacts on individual autonomy and rights. This level of data collection about individuals and data sharing by the state can turn sinister, and rapidly slide into mass surveillance of a populace, the vast majority of whom are under no suspicion.

Like private sector data collection, this carries risks to autonomy and the risk of manipulating behaviour. This is not only through active decisions by government to 'nudge' behaviour, but also via more generalised chilling effects: people who know they are being watched may feel compelled to change their behaviour.

Consider government reliance on data-driven predictive analysis to make decisions. The more data is linked, the more fine-grained analyses, decisions, and even interventions governments can make, in relation to particular identified individuals or, more likely, in relation to non-identified subgroups of people based on their shared characteristics.

There are genuinely difficult questions about the extent to which we want governments making predictions about the life trajectories or likely activities of individuals, let alone intervening, by offering or refusing access to services based on those predictions.

The intention can be benign: for example, to offer tailored support to improve people's opportunities in life. But even that kind of benign intention involves not offering additional support to some people who need it; predictions are not 100% correct, and incorrect predictions are not always randomly distributed. In one US instance, black patients missed out on additional medical resources because they were disadvantaged by an algorithm that used past spending (which historically had been lower for black patients) to predict future need (Obermeyer et al. 2019).

Even if we assume predictions are accurate, when official organs make a prediction about us, rather than a judgement based on past performance, it is far from clear whether we are being treated as individuals and according to our own merits. And this raises questions of justice where it affects a person's opportunities in life.

When final school exams were cancelled in the UK due to the COVID-19 pandemic, students were initially assigned A-level grades based on the operation of an algorithm that sought to predict their performance based on past performance of their school. Many results were downgraded from school-assigned grades, disproportionately in the state-funded educational sector.

Public uproar caused authorities to retreat and reinstate teacher assessments. In part, the controversy concerned the distribution of the impacts, with schools with lower socio-economic student populations more adversely affected. But there was also a sense that there is something wrong with determining where a student goes to university, or what they can study, based not on how they have actually performed but on how a computer predicted they would, and based significantly on other past students' performances (even if the underlying purpose was to predict likely success at university studies).

The latter feels like science fiction predestination, and not seeing the individual for themselves. The idea of predicting performance rather than relying on actual performance has been discussed in the Australian university sector as well, even before questions were raised about examinations through the COVID-19 pandemic.

Questions of justice are even more acute if predictive analysis is used in law enforcement. Data is an essential tool in modern law enforcement. But the use of prediction in shaping policing activities (such as where police cars patrol), treatment in the criminal justice system (for example, whether bail is approved or denied) or administrative treatment (such as who is offered a payment plan) can have significant impacts on individual freedom, leading to certain kinds of people being repeatedly targeted or ignored.

It can also lead to perverse outcomes: as where people change their behaviour (who they socialise with, what they say on social media) in order to optimise their treatment by government or private sector organisations. Not only can this undermine the basis for prediction (as correlations will change), it can also have perverse broader social effects. For example, if government were to prioritise fine collection from those who it predicts will pay promptly (based on credit records), this would create a perverse incentive to pay bills after they fall due.

The COVID-19 pandemic and its associated social/economic crisis has accelerated trends in data collection and use. As a result of COVID-19, both governments and private sector actors have rapidly accelerated data sharing and linkage; the application of data analytics to highly sensitive health data; and innovations in all kinds of surveillance to increase awareness of where people are, how they are moving around, and who they are coming into contact with – on an aggregate and individual basis.

COVID-19 has also highlighted the potential for the intermingling of private and public sector data collection and use. This raises its own serious challenges. It is by no means clear that members of the public, before the pandemic at least, were comfortable with public sector use of private sector data, or (particularly) vice versa (Goggin et al. 2017). But the pandemic saw, for example, public authorities using (aggregated and anonymised) privately held mobile phone data to monitor whether people generally were complying with restrictions on movement.

In some overseas cases, this extended to targeting specific quarantined individuals, and in 2021 we saw the first discussion of the potential for use of technology for similar purposes in Australia. Governments also moved to encourage the collection of new kinds of data, via apps designed to log close contacts between people, such as the Australian COVIDSafe app, and more importantly, the near-compulsory use of QR code check-in apps, with data collected by government.

COVID-19 illustrates the underlying potential for shifts in our intuitions about what is okay, and what is not – and that we can make rapid policy changes. Whatever its flaws or effectiveness, it is striking to observe that millions of Australians proved potentially willing to reveal who they've spent time with by downloading COVIDSafe, and millions

more have provided a detailed map of their locations via QR code check-ins.

It is notable too that the Australian government showed itself uncharacteristically modest by amending the *Privacy Act 1988* to restrict how COVIDSafe data might be used: protections beyond that granted to other datasets in Australia. On the other hand, similar protections were not enacted for QR check-in data, leading to several instances where state police forces have accessed the data for investigations unrelated to COVID-19.

All this makes it clear that it is a challenging – but critical – time to be thinking about data and how it is used across public and private sectors, and how much of that use is legitimate. Questions are still open as to whether, and which of, these pandemic-related shifts will and should become permanent; and whether privacy protections will be extended or not.

As a society we need to come to some kind of consensus about what is allowed across the scenarios above and more, and we will need to keep reviewing and renewing that consensus over time. We are unlikely to solve the legislative and design problems we face unless we can come to some agreements around what is legitimate, and what is not, in these scenarios and others, and around the underlying principles that should inform how we respond to them.

Here, we are not offering a final answer to any of these points. When we talk about ‘coming to’ agreement, we are acknowledging that societal consensus does not currently exist; that it’s something we’ll need to actively work on as a society and not just as experts (as to which we have some ideas, elaborated in Chapter 5). Importantly, deciding how far we are prepared to go in each of these scenarios will be an ongoing process, not something to ‘set and forget’. As technology and uses of data change, and as we better understand the implications of those changes, we need to constantly rethink what we want law to encourage, discourage, and prevent.

Locating the problem in existing legal frameworks

The first reference point for addressing the scenarios above is the existing law. Australia has a patchwork of laws governing data processing and data-driven targeting and decision-making. Although a wide range of other laws could potentially impact the use of data, as the most immediate and relevant, we focus here on privacy law, anti-discrimination law, tort law, and emerging attempts to outline ethical (or possibly legal) principles for AI. We address the extent to which these, or amended versions of these, might usefully resolve the kinds of issues discussed above.

As this section will show, existing laws do not cover the range of problematic situations described above. Some issues can be solved by tweaking law – for example, enhancing privacy law through a shift towards the European approach exemplified in the GDPR. But, in other cases, the goals, focus and provisions of existing laws are not well directed towards the problems being encountered.

Privacy law's notice and consent model fails to recognise absence of real choice; anti-discrimination law can be difficult to invoke in the context of opaque machine learning algorithms and fails to capture broader harms; and regulation targeted at particular technologies (such as automation or AI) is insufficiently flexible. While we do not analyse all relevant legal regimes here (such as consumer law),²² they are similarly confined to their own purposes, which fail to capture the full scale of the problem.

Privacy

The obvious legal frame for tackling our assorted data problems is privacy: it is our first port of call when looking for 'laws that control data'. But what do we mean by privacy?

Privacy is a multihued concept with many definitions. It is variously described as our right to protection against intrusion into seclusion; protection for having 'breathing room' or a 'safe haven' to be and develop ourselves without scrutiny (Cohen 2019); our right to decide to whom we reveal information about ourselves, and for what purposes and in what contexts; and a demand for respect for context-specific norms for information gathering and use (Nissenbaum 2004).

We assert our 'privacy' to push back against everything from neighbours looking over the fence; exes circulating intimate photos or communications; service providers and employers seeking excessive or intrusive information; and the government monitoring its citizens. In the US, the right to privacy even protects liberties such as the use of contraception and parental rights over choices relating to their children.

These various rights, concerns and interests all justify privacy protection, but we're concerned here with privacy *law* in Australia: chiefly as legislated in Australia via the *Privacy Act 1988* (Cth) as well as other piecemeal Commonwealth and state legislation. This existing privacy legislation does not seek to provide a general protection for privacy, but rather focuses on control over personal data in particular contexts. For the sake of comprehensibility, we focus here on basic privacy laws: not new specialised schemes like the Consumer Data Right, or the specialised rules in national security and law enforcement, or the tailored regimes for health data or COVIDSafe app data.

Many have commented that Australia's current privacy laws are inadequate, especially in the context of the 2020 to 2022 review of the Privacy Act. Some of these problems are well known and easily addressed – such as the lack of any private cause of action when privacy law is breached. In our contribution, we want to pick apart the more fundamental reasons why the multiple issues we face today won't be addressed by incremental amendments to privacy law. Even more wholesale change along the lines modelled in Europe's GDPR may not be sufficient. A key part of the problem is that certain assumptions were made by policymakers back in the 1980s when they were writing the principles underlying current privacy legislation, although they probably did not even realise they were making them. Those assumptions no longer hold.

²² On the limitations of consumer law, see Manwaring (2018).

Today's scenarios highlight the difference between *identification* and *personalisation*. Back when current laws were being written, we assumed that data collection and use only threatened our vital interests or personal autonomy if it could be linked to our name (or some unique identifier), and used against us or revealed, without our consent. Reflecting this assumption, privacy legislation only applies to information where the person is, or can reasonably be, identified (*Privacy Act 1988*, s 6), and conversely, anonymisation or de-identification of data is presented as a solution to privacy concerns. But not only is complete de-identification inconsistent with data utility (Ohm 2010), personalisation is not the same as identification, and so anonymity is no longer the protection it was (Barocas and Nissenbaum 2014).

In a world with big data analytics, businesses and governments can target individuals with personalised information, services or offerings without ever knowing their identity. And identity is no longer the key to finding out more about a person. A business that knows just enough about me, and a lot more about other people, can use data analytics to make inferences to fill in the facts it doesn't know. This is why Anna Johnston has talked about the need for privacy legislation to address *individuation* – the ability to 'single out' a person for tracking, profiling, targeting, contacting or a decision, whether or not they can be identified (Johnston 2020).

Our privacy legislation is also structured around a distinction between primary and secondary uses of data. It posits that data is collected for some purpose reasonably necessary or directly related to the collector's activities, and then imposes legal limits on other, 'secondary' purposes.

This rather static conceptual distinction breaks down in the current dynamic environment. What are the 'purposes' or functions of an all-encompassing behemoth like Amazon? Or a large, complex online platform like Facebook, operating a two-sided market serving both non-paying users and paying clients seeking to show material to heavily targeted subsets of those users, and others – including governments – seeking broader insights about people and their behaviour?

As commentators such as Cohen (2019) have noted, the whole business model of platforms is dependent on extensive data collection and use – for a multiplicity of purposes, presently known and unknown.

In this context, the OAIC pleading, in a current case brought against Facebook, that the 'primary purpose' for which Facebook collected information was to enable people to build an online social network with other Facebook users seems oddly quaint, even if it is appropriately narrow, as you would expect from a privacy regulator. Put simply, today's data analytics and machine learning break the 'limited purpose' model, because their whole orientation is to use large collections of data to be analysed for *unexpected* connections or insights.

And this brings us to the well-known challenges with privacy legislation's notice and consent model. Others have discussed the many problems with consent, and

the ACCC has explored how we might improve the model in some depth (Australian Competition and Consumer Commission 2019).

We would add that there is reason to think that a model based purely or mostly on individual control and consent will not address all of the potential concerns and harms that arise out of today's scenarios however much we improve it, or make consent disaggregated or opt-in, or require more readable privacy policies.

In the context of today's complex information economy, a model based on consent faces fundamental paradoxes. Given an infinite cloud of senders and recipients of data, ongoing and dynamic changes to the functions and purposes of data collectors such as Amazon or Google, and the rapidly developing capacity of data analytics, it is not possible to communicate to people what they are consenting to in ways that are both meaningful and simple.

A simple and brief explanation is necessarily abstract ('we will use this data to choose what advertisements to show you'), but the real devil is in the details ('we will use data about what you just said on social media to decide how you are feeling, and target ads for products we know people buy when they're feeling that way, as well as stories that will keep you on the platform longer'; or 'we will infer your political leanings from what charitable causes you support and what news sources you read, and show you even more of that, reducing variety in your media diet and guaranteeing you won't have much in common with your relatives who have different political leanings').

Consent also assumes a real *choice*. Despite the fact that many services such as job listing platforms are virtually essential, many platforms offer 'take it or leave it' privacy settings, and people lack time to absorb all of the often complex information in privacy policies (Leonard 2020).

This is even before we address the fact that genuinely essential services – such as electricity – increasingly come with data collection. To make matters worse, many sites use a variety of techniques to manipulate users into settings that are less protective of privacy (Pardes 2020). And as pointed out above, given that a key purpose of modern machine learning is to uncover unexpected connections and insights, consent given ahead of the fact will necessarily be incomplete.

Even if we could address these paradoxes, there are situations where notice and consent is just not the right frame for thinking about how we address problematic uses of data, because not everything we value about either data *or* privacy is purely individual.

There are also collective interests, and if we want to recognise and promote collective goals, then we can't rely on individualised consent (Viljoen 2020). For example, if we rely on consent as the model for making data available for social and medical research, we put at risk the considerable potential collective benefits (Australian Government Department of the Prime Minister and Cabinet 2019).

We also cannot meaningfully give or withhold consent to government data-gathering, which is necessary to make political and social systems function. This became even

more acute in the COVID-19 pandemic, where personal location and health status information was used to protect everyone.

More generally, privacy is a collective good: *my* privacy (what people can know or infer about *me*) depends on *your* actions (what information *you* choose to disclose *about you*). Especially in the context of big data analytics, if some people volunteer their information, then analysis can fill in the gaps for the people who withhold consent.

As Barocas and Nissenbaum (2014) have noted, 'once a critical threshold has been reached, data collectors can rely on more easily observable information to situate all individuals according to these patterns, rendering irrelevant whether or not those individuals have consented to allowing access to the critical information in question. Withholding consent will make no difference to how they are treated'.

Our traditional legislated privacy rights are important. Governments and companies *should* have to think about the purposes for which data is collected, be transparent about how data is collected and used, seek agreement, keep information secure from unauthorised disclosure or use, and disclose and remediate data breaches.

Current legislated privacy rights provide a baseline and are built into technical and management systems and a network of international treaties and similar laws in other countries (Greenleaf 2019). We should hold on to them and encourage their implementation through privacy by design and privacy-enhancing technologies focused on protecting data security and anonymity.

As discussed by many in the 2020 to 2022 review of Australia's Privacy Act, we should update privacy law, perhaps with aspects of the EU GDPR that offer useful incremental reform, such as further restricting secondary uses of data and enhancing transparency and consent requirements, as well as creating a stronger enforcement and penalty regime and obligations to build privacy into product and service design. What our discussion tries to suggest, however, is that to address our data dilemmas, we also need to look beyond traditional privacy frames.

Tort

In some jurisdictions, such as the UK and Canada, there is a claim in tort for infringement of privacy in personal information (Trakman, Walters and Zeller 2019). Article 82 of the GDPR also provides for compensation in the event of damage that results from an infringement.

A similar tort or statutory right in Australia, as recommended by the Australian Law Reform Commission (Australian Law Reform Commission 2014), would overcome one limitation of privacy law, namely the lack of a right of private action. It could also extend legal rights to prevent unreasonable intrusions upon individual seclusion, and more effectively provide remedy where strangers collect information on us, outside of the ordinary commercial and government transactions that are addressed through privacy legislation, for example, where your neighbour sets up a camera that films your backyard.

However, there are practical difficulties with relying on individual plaintiffs to bring proceedings for damages. Most obviously, there is the expense of doing so. Equally problematic is the difficulty of proving substantial quantifiable harm (Trakman, Walters, and Zeller 2019). Further, many of the general concerns about privacy law would remain, including the fuzzy boundaries of personal information, the individual focus of the remedy, and the ease of obtaining consent for complex, intertwining data practices.

Also, there are there no current signs in Australia of any move towards broader use of tort to protect against harmful or careless use of linked data in ways that harm data subjects, at least as it concerns government use of data about individuals.

In *Prygodicz v Commonwealth of Australia (No 2)* [2021] FCA 634, one of the two broad claims put forth on behalf of individuals subjected to so-called 'robodebt' claims was negligence, specifically that the Commonwealth breached its duty to exercise reasonable care in the performance of Commonwealth-controlled functions under the *Social Security Act 1991* (Cth) to avoid foreseeable economic loss to the applicants.

In his judgement accepting the negotiated settlement, the judge considered this claim weak, doubting the applicants could establish the novel duty of care alleged. The fact that the pleaded case for a duty of care was considered weak does not provide much hope for the use of tort law to discipline shoddy government activities using data – especially when set against an incident otherwise described by the judge as 'a shameful chapter in the administration of the Commonwealth social security system and a massive failure of public administration'. On the other hand, proposed class actions arising from the massive data breaches that occurred in 2022 provide an opportunity to consider private sector data practices and tort.

Discrimination

What many fear about automated processing of their data is not that there will be intrusion on their seclusion, nor that personal information will leak out: gender and racial identity is often public or easily available. The fear rather is that we will be treated differently from others for arbitrary reasons. Examples of this were laced through our introduction.

This concern is not confined to individuals. Each of the above examples we gave tends to perpetuate stereotypes in society more broadly. Directing certain job ads away from women makes it more difficult for women to see themselves as engineers; targeting rental advertisements only at societal subgroups makes for less diverse neighbourhoods.

There are further examples where the harm is *primarily* societal rather than individual – as where Google image searches for 'baby' show a disproportionately high number of white babies, or where predictive policing software sends more law enforcement officers to racialised neighbourhoods. In contexts like Google image searches and predictive policing, the data processed may not be about a reasonably identifiable individual, which places it outside data protection laws.

Anti-discrimination law can play a role to prevent some of the more egregious harms of targeting and profiling. The precise laws differ by jurisdiction, but Commonwealth law in Australia prohibits discrimination on the basis of sex and race in some contexts.

Anti-discrimination law generally distinguishes between direct discrimination (sometimes known as disparate treatment) and indirect discrimination (sometimes known as disparate impact) (Gaze and Smith 2017, 22–23). This difference is between a situation where an attribute is used as a basis for a decision or where a decision is 'by reason of' an attribute (direct), and a situation where, despite the attribute not being used directly, there is a negative impact on people with a particular attribute (indirect).

Direct discrimination is easy to avoid in the context of automated data-driven inferencing, particularly where protected attributes are not used as variables by the system. The COMPAS tool used in risk assessment in the criminal justice system in *inter alia* Wisconsin in the US did not, apparently, use race as an explicit variable (Angwin et al. 2016).

Even where discrimination is direct because it is 'by reason of' a protected attribute, proof is often difficult. The algorithm that filtered out female applicants for a position at Amazon was not programmed to use gender as a variable relevant to its analysis. Rather, it learnt that the presence of particular words in a resume (for example, reference to a 'women's' college or sports team) were less aligned with the resumes of those already employed and performing well in a particular team.

Thus, although in such a case one could say that the algorithm is making decisions, at least to some extent, 'by reason of' the presence of female words in a CV, *proving* this requires an opportunity to analyse the algorithm. Because it was not programmed in directly but was rather learnt from training data, inspection of software specifications would be insufficient.

Indirect discrimination has the advantage that it can be measured even though the algorithm itself is a black box. In the hiring algorithm example, it could be shown that the filtering condition had the effect of disadvantaging women. However, indirect discrimination is subject to more legislative exceptions than direct discrimination. An employer can argue, for example, that it receives so many applications that reliance on a machine learning filtering process is reasonable. The success of the argument would depend on the nature and extent of disadvantage, the feasibility of overcoming or mitigating the disadvantage, and whether the disadvantage is proportionate to the result sought (*Sex Discrimination Act 1984*).

Coupling the fact that direct discrimination has fewer defences than indirect discrimination together with privacy law's rules limiting the gathering or holding sensitive information tends to encourage data processors to omit protected attributes as variables in the analysis. This is not an optimal way of ensuring that an algorithm treats people fairly, as there are many variables that will correlate with protected variables, and removing variables of interest from the analysis makes disparate impact hard to measure and correct.

Discrimination law is also of little assistance in addressing broader harms, such as racially skewed output in a search for images of babies. There is no individual aggrieved by such searches, even though there is societal-level harm as people generally are socialised to think of babies as primarily white. These harms propagate out through a range of channels: white babies turn up in the searches, white babies are disproportionately selected for presentations and brochures, the diversity of the world we actually live in is obscured.

Discrimination law is also unhelpful outside pre-identified protected attributes. There is no discrimination law that protects against pricing algorithms charging Safari users more than Chrome users; no category that protects against algorithms for assigning A-level scores that downgrade the results for more students in government-funded than private sector schools.

More problematically, discrimination law does not help ameliorate targeted political advertising that divides a population by psychological profile and targets messages according to susceptibility to persuasion along different dimensions. Being susceptible to certain kinds of arguments or emotional triggers is not a characteristic protected by discrimination law.

It is unlikely that these limitations of discrimination law could be solved through revision of the existing law. The purpose of discrimination law is protecting individuals from disadvantageous treatment based on particular predefined attributes.

While it may be possible to rethink the distinction between direct and indirect discrimination, particularly where decision-making processes are intentionally data-driven, as in the context of machine learning, this would be a radical break. Were such a rethinking to occur, it would be important to engage with data scientists around the metrics that such systems might satisfy, at least with respect to differential treatment of people according to categories such as race and sex. In such a way, discrimination law can be improved, but even then, it cannot prevent all the ways in which people are treated differently based on data-driven inference.

Governance of AI

Many have identified that the data dilemmas we face are associated with AI or automated processing. This has led to a rush to promulgate ethical principles to govern AI and automation (under various definitional guises) with reference to terms reminiscent of privacy and discrimination law (seen in the demand for 'fair' AI), but also extending to accountability, transparency, explainability, sustainability, robustness and resilience, and beneficence.

In Australia we have seen the Australian Government formulate a list of AI Ethics Principles (Australian Government Department of Industry, Innovation and Science 2019). These overlap with many principles produced in similar documents (Fjeld et al. 2020) and include the need to respect human rights and human autonomy, the importance of fairness and inclusivity, the need for reliability and safety, and the requirement for transparency and explainability.

None of the principles are wrong, although all of them would seem to be equally important whether or not AI is involved. It is difficult, for example, to understand why the government is focused on whether AI systems benefit the environment, as opposed to introducing better broadly applicable environmental laws and regulations (or even generally applicable ethics principles).

Along similar lines, the Australian Human Rights Commission (AHRC 2019) has made a number of proposals concerning 'AI-informed decision-making', such as legislation that would require that *inter alia*:

- individuals be notified where AI is materially used in making an administrative decision (recommendation 3)
- individuals be notified when a corporation or legal person materially uses AI in a decision-making process that affects the legal or similarly significant rights of an individual (recommendation 10)
- reasons be generated or a technical explanation given to those affected by administrative decisions that use automation or AI (recommendation 5).

In addition, recommendation 16 suggests a human rights approach to government procurement of products and services that use AI, with relevant laws, policies and guidance amended to require the protection of human rights in the design and development of any AI-informed decision-making tool procured by government.

Such recommendations are laudable, but potentially too narrow. Why should such a rule be limited based on the technology involved? In other words, if the government is procuring goods or services not involving AI, are we no longer concerned about the impact on human rights? Do we want human rights protection and promotion because of and only in circumstances of technical mediation (AI/automated processing), or more generally?

There is a similar flaw in the new proposed EU regulation on AI.²³ For example, Article 5(1)(a) would prohibit:

The placing on the market, putting into service or use of an AI system that deploys subliminal techniques beyond a person's consciousness in order to materially distort a person's behaviour in a manner that causes or is likely to cause that person or another person physical or psychological harm.

But what if one were to remove the term 'AI', so that the regulation would prohibit the use of *any system* that had that effect. What is it about AI that suggests that the same harm requires a different response? The narrowing of the definition since the 2021 draft suggests that the scope of what precisely ought to be regulated is still subject to debate. But the point here is that there is *no* definition of AI for which human rights or psychological harms are uniquely important.

23 Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM(2021) 206 final (21 April 2021). Updates are still being proposed and incorporated since that draft.

The same point about the limitations of a technologically specific approach (even a broad one that covers all AI) applies to other regulation.

One example is the regulation of decisions 'based solely on automated processing' in Article 22 of the GDPR. That article creates a right for data subjects 'not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.'

There are broad exceptions, which include automated processing necessary for entering into or performance of a contract between the data subject and a data controller, processing authorised by other laws that provide suitable safeguards, and where the data subject gives their explicit consent. Where those exceptions apply, the data controller must still implement suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, involving at least the right to obtain human intervention. Additional rights that apply in some circumstances offer 'meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing' for the data subject, see GDPR articles 13(2)(f), 14(2)(g) and 15(1)(h).

The GDPR approach could manage some of the concerns about personalisation that denies individuals opportunities, and discrimination based on black-boxed data-driven algorithms. However, there are important limitations. A crucial one is the exception where the data subject has consented, which is often freely given in the context of click-wrap online terms and conditions (Mendoza and Bygrave 2017).

But the most significant limitation relates to the degree of automation involved – *all* protections are linked to decisions based *solely* on automated processing. Where there is human intervention in the process, arguably excluding pure rubber stamping, the fact that the automated output is highly correlated with the final decision will not be enough to trigger the various protections (Veale and Edwards 2018).

All of this points to a key question encountered in Australian approaches to AI ethics: to what extent are our concerns about automation or AI as opposed to the fact that data – or more data, or different data, or data about us we might consider irrelevant to the matter at hand – is being used to make decisions that affect us?

If the data processing were manual rather than automated (for example, if humans made decisions based on statistical modelling geared at predicting the behaviour of data subjects), would that address the core concern? Consider again the examples of women being shown different advertisements for jobs or people of colour being shown different advertisements for housing – would it help to have a human involved in the process if they still relied on data-driven inference (perhaps looking up multidimensional tables that purportedly gave likelihoods of interest in jobs or suburbs)?

A similar point can be made about the frequent demands that AI should be explainable. The challenge presented by machine learning (whether presented as automated processing, AI, or some other term) is less about *whether* internal

reasoning is made clear or rational, and more about *how* this can be done (at least to the same level as human decisions).

Some machine learning algorithms are too complex, basing inferences on complex patterns and relationships in data that are difficult for even expert humans to disentangle. Where particular categories of decisions must be explainable or demonstrably rational, those algorithms are simply the wrong tool for the job.

We can solve the problem by setting technology-independent requirements for certain categories of decisions (for example, administrative decisions must be accompanied by reasons, employment decisions must be demonstrably independent of race and gender, and sentencing or insurance decisions must be based on a limited range of variables). Such requirements may drive useful research in computer science, for example, algorithms that verify that other algorithms satisfy particular criteria or generate reliable explanations for their behaviour. They may also drive useful standardisation as to the means through which algorithms can meet particular common criteria (potentially with certification).

The core question remains – what is it that we are concerned about when we discuss these kinds of problems? Why do we want explainability or rationality in how we are treated, and in what circumstances should we have a right to demand it? Are we really worried about automation and technological sophistication, or is there a deeper underlying concern of which particular technologies are merely a modern manifestation?

Technology-specific regulation seems too narrow a solution to the dilemmas made manifest by new data practices. The problem is less automated processing or AI than a mode of decision-making where people are placed in categories, so that assumptions can be made based on the actions of other people who are sufficiently 'like' them. Especially when detrimental consequences follow. And most especially when the basis of categorisation is unable to be explained, or contested. Understanding why this is a problem, and locating starting points that address it, is the subject of our next section.

Other starting places

Each of these existing frameworks, as it currently exists, offers only a very partial answer to our current data dilemmas. In this section we discuss some alternative starting points. We confine our discussion here to two principles that we think people in Australia would generally consider fundamental to our system of governance: the rule of law, and respect for human dignity and autonomy. We have chosen these in part because these principles are already deeply embedded in Australia's political and legal system, and generally understood to be important by policymakers and the public alike. We could have discussed others – trust, for example, is another important element of our system of governance; a human rights frame is another, adopted by the AHRC.

Rule of law

One of the primary concerns about the way governments and large technology companies use data-driven inference to direct the lives of individuals is the potential for abuse of power. Power is evident in the way that social media and search engines direct our attention to content in ways that align more with corporate objectives than individual curiosity and serendipity.

It is also evident in the manner in which important decisions affecting us (determining entitlements, making hiring decisions, adjusting prices, making offers) are made by systems we do not understand, and which we have limited or no ability or opportunity to challenge.

Returning to the data problems raised earlier, the potential abuses of a national scheme of digital identity based on facial recognition, the use of Facebook's powers of influence people to manipulate elections, and racialised impacts of predictive policing are all frightening because they involve potentially arbitrary exercises of power.

These exercises of power are arbitrary either because they involve judgements about us but based not on what we've done but on what other people do or have done; or because they involve seeking to shape our behaviour in ways hidden from us towards ends that are not our own in the absence of rational or moral argument.

Where abuse of power has been a concern historically, the rule of law has arguably been the most successful solution. The rule of law is a public law concept whose purpose is to temper exercises of power (Krygier 2019). The rule of law comes into play in a context where one actor exercises significant power over others: it is what Australia (and many other countries) rely on to ensure that such power is not exercised arbitrarily.

The classic example is government, but the same principles ought to apply for corporations such as Facebook with significant control over their users' online social networking (Krygier 2011, 88–89). Here, the focus is not on the technological means through which power is exercised, but rather basic minimum standards required of powerful actors when exercising their power over others. Those minimum standards lie beyond other rights and interests – individuals' consent and corporate interests in preserving trade secrets do not override them.

The manner in which the rule of law tempers power is theoretically contingent, but in Australia (and elsewhere) is commonly tied to fundamental values such as accountability and transparency, predictability and consistency, and equality before the law. Interestingly, these overlap significantly with ethical principles that are commonly said to apply to AI.

The primary difference is the standing of each – we can ask corporations and governments to act ethically, whereas we should *demand* that at least governments observe the rule of law, and label government actions illegitimate where they are not consistent with the rule of law.

A rule of law lens is particularly useful where the concern is about our susceptibility to arbitrary decisions that cannot, in practice, be avoided, so that consent is meaningless. Consider the use of COMPAS in bail, sentencing, parole or incarceration decisions, or an algorithm used to assign final grades to school students. Rather than limiting arguments to anti-discrimination law, the rule of law value of equal treatment under the law arguably proscribes the treatment of all arbitrary variables (or variables simply thrown in a bag for identification of statistical correlations) (Zalnierute, Bennett Moses and Williams 2019). A rule of law argument undermines the idea of data-driven decision-making, imposed by the powerful, in ways that are designed to discriminate based on potentially arbitrary criteria.

The rule of law lens does, however, have important limitations. Most crucially, its meaning is often contested, so that while in theory it enjoys broad international acceptance, its diverse interpretations mean it does not always work in the same way, or serve the same function, around the world.

Even in Australia, it can be argued that differential treatment relying on data-driven inferences can be rationally justified. So what we are describing here is just one possible interpretation of the rule of law, one designed to ensure a measure of protection for individuals, perhaps at the expense of rational scientific management of populations. The rule of law is also minimalist and fails to capture the broader agenda associated with human rights and human dignity.

Human dignity and autonomy

Another way of looking at questions of power and subordination common to some of the data problems described above is through a lens focused on human dignity and autonomy. Both concepts lie at the heart of international human rights instruments, and they can be powerful principles for thinking through the impact of large-scale data collection, predictive analysis and automated decision-making.

Like the rule of law, people argue about what it means to have or respect autonomy or human dignity. Still, we can say that these concepts have some core meaning. By 'autonomy' we mean the idea that people have final authority to control their own lives. It reflects the assumption, foundational to the Western liberal tradition, that people have the capacity to develop and act on higher order plans of action: to make their own decisions about the life they want to live, and that providing space for people to do so is important (Richards 1981).

The core to the idea of 'human dignity' as we're using it here encompasses the principles (McCrudden 2008):

- that every human being possesses intrinsic worth, simply by reason of being human
- that this intrinsic worth should be respected by others, and reflected in how people are treated
- that the state exists for the sake of individual human beings, not vice versa.

In part this means recognising that people are to be treated as individuals in their own right, and are not mere objects, or things, to be used or manipulated in pursuit of others' ends, whether those 'others' are powerful private actors, or the state.

Seeing our data problems through this twin lens casts them in a different light, and finding ways to embody concepts of autonomy and human dignity could help fill conceptual gaps in current rules. Privacy, both in the broad sense of 'seclusion against intrusion' reflected in Article 17 of the *International Covenant on Civil and Political Rights* (ICCPR) and the narrower sense of control over data collection reflected in Australia's *Privacy Act 1988*, is a means of protecting autonomy. As Cannataci (1987) has noted:

Shorn of the cloak of privacy that protects him, an individual becomes transparent and therefore manipulable. A manipulable individual is at the mercy of those who control the information held about him, and his freedom, which is often relative at best, shrinks in direct proportion to the extent of the nature of the options and alternatives which are left open to him by those who control the information.

Where data-driven predictive analysis is used to dynamically tailor some systems specifically to manipulate our behaviour to achieve some goal of the system, this compromises our ability to make our own decisions consistent with our own assessment of our interests and goals.

Systems imposing this kind of harm would include, for example, deliberately confusing or manipulative commercial techniques designed to exploit human weaknesses, or user interfaces deliberately designed to make it harder for people to exercise their preferences, for example, for data use.

Clearly there are questions of degree here: arguably all modern advertising is deliberately manipulative in some way: but recent research suggests ways to distinguish between ordinary advertising and the kinds of 'dark patterns' that should be treated as too aggressive or manipulative (Luguri and Strahilevitz 2019).

More fundamentally, the lens of human autonomy and dignity assists in understanding why data-driven predictive analytics is troubling even if well-intentioned. Systems that seek to predict peoples' behaviour or life trajectories based on their similarities with other people, especially those that *intervene* on the basis of such predictions, potentially undermine both autonomy and human dignity.

The fear is that technologies that use large-scale data analysis to categorise and characterise individuals, and intervene to alter their behaviour, risk 'subordinat[ing] considerations of human well-being and human self-determination' to the priorities of others, whether commercial interests, government, or both (Cohen 2019).

Forecasting how people will act based on their similarities with others – and intervening to affect those choices before they have been made – necessarily involves treating the person not as an individual but as a kind of object. It is also inconsistent

with respect for human dignity to treat people differently based on personal features, characteristics, or circumstances that are unrelated to their needs, capacities or merits.

So, for example, concluding on the basis of a quantitative analysis that a person falls into a higher risk category, and denying them some benefit (parole, or employment, or an educational opportunity) because their parents divorced when they were young, or because they wear size 10 shoes, is inconsistent with a respect for human dignity, even if the maths is accurate.

Like the other frames we've discussed, a focus on human dignity and autonomy is an incomplete way to think about today's data problems. Not least, such a frame encourages us to think individualistically about the impact of technology, where many issues arising today affect us collectively.

Individual autonomy cannot be the sole lens, for example, for thinking about the availability of data for medical research, or we might sacrifice the good of medical research in order to avoid treating people as objects. It would be wrong to treat these principles as absolutes; rather, they illuminate justifiable concerns, and require that affronts to these principles be carefully justified, proportionate, limited and controlled.

Thus medical research that uses individual data can be justified where proportional and where it doesn't harm the individual. Still, these principles give reasons why data-driven predictive analytics based on analysis of 'people like me' are a troubling affront: reasons that are absent from traditional privacy, tort, and anti-discrimination laws, and more satisfying and complete than the attempt to write technology-specific regulation.

A lens that includes consideration of human dignity and autonomy may also provide an easier foundation for imposing controls and obligations on the activities of the private sector. On some interpretations – which we do not agree with – the 'rule of law' is concerned only with the control of *government* power. Even those who argue that the rule of law is no concern of the private sector, however, will find it harder to argue that private actors have no obligation whatsoever to respect human dignity and autonomy. 'I can manipulate and treat people as objects for profit without constraint' is a very unattractive position to take.

Making it real

If we are right, principles such as the rule of law, and respect for human dignity and autonomy can provide a useful frame for understanding why developments in the mass collection and use of data, predictive analytics and automated decision-making are a problem for Australians.

These principles are already deeply embedded in Australia's political and legal system, and generally understood to be important by policymakers and the public alike. But we still need to answer the question: how do we make those principles a concrete part of the way we address current data problems, as a society?

Principles like the rule of law, or respect for individual dignity and autonomy, do not lend themselves to easy translation into concrete legal obligations. There is no 'international human right to the rule of law', and respect for human dignity is explicitly included in the ICCPR only partially, via Article 10, which talks about the treatment of people deprived of liberty. As principles, they operate at a higher level than legislation like the Privacy Act. But that does not make them any less important. We can, and should, look for ways to use them to address our current data dilemmas.

We would suggest that these principles are important at three levels:

- in legislative reform
- in the discourse around privacy, data and artificial intelligence to counter what can otherwise be a narrow focus on existing privacy laws
- as a consideration in the design of systems.

In other words, our laws, discourse and systems should reflect the importance we place on these fundamental values.

Legislative reform

If we are going to reform privacy law, which is the subject of discussion in Australia at the time of writing, then respect for the rule of law (and the prevention of the arbitrary exercise of power) and respect for individual dignity and autonomy ought to inform any changes. We should have laws that ensure accountability of government decision-making, facilitate human autonomy and dignity, and protect against arbitrary exercises of power by governments and corporations. This will require, *inter alia*, reforms to privacy law and discrimination law but much else besides.

These principles also need to be much more actively thought about in other data-related legislative reforms and implementation. Consider, for example, the *Data Availability and Transparency Act 2022* (Cth) and its ongoing implementation, and similar laws being written in the states to govern linking and sharing publicly held data about people.

Such laws are often discussed as if protecting privacy (for example, through de-identification) is sufficient for a legal framework that protects people's interests, and that any concerns or risks beyond that can be dealt with through ethical frameworks. But if we take seriously the demand, from rule of law principles, that arbitrary exercises of power be controlled or prevented, and the demand, as a matter of respect for individual human dignity, that people be treated as ends in themselves rather than as objects for the achievement of others' goals, then we might start to take seriously – and legislatively control – the potential harms arising from data linkage, the inclusion of facially irrelevant data in analyses, and the ability to make predictions about people based on their characteristics rather than their identity.

In addition, the rules themselves need to be clear and transparent, so that they are easily understood by data subjects and by corporations, law enforcement agencies and governments. This understandability of law is as essential for the rule of law as the content of laws.

Laws should be reviewed not only for their substance, but for clarity, consistency of terminology, intersections across a complex web of legislation, and ease of navigation by those impacted. This is not currently the case; for example, there are over 50 words and phrases used in legislation concerning powers and responsibilities for data to identify the entity that has those powers or responsibilities for particular data (Bennett Moses 2020). We can create more integrated and comprehensible legal frameworks.

Discourse

We need to rethink how we talk about our current data dilemmas. Currently, in Australia, there is a tendency to frame data collection and use as a privacy issue, and to talk about privacy as either 'dead' or as needing to give way to serve other interests (such as security and health).

The fact that people click 'I agree' to privacy policies they lack time to read or the fact that they use social media to connect with friends is taken to be a sign that they are not interested in privacy at all. But this is not necessarily the case: research in the past has shown that people hold nuanced views about data use (Goggin et al. 2017).

It is certainly true that concerns about the collection and use of data are not on most lists of what matters most to Australian citizens. The Australian Broadcasting Corporation's 2019 Vote Compass found voters split on whether the economy or the environment was the most important issue, and data practices were not even on the list of options in the survey (Hanrahan 2019).

While data practices intersect with some important issues (such as government accountability and inequality), there are relatively few voters who would be persuaded to change their vote based on privacy law reform. Nevertheless, community views are an important part of the story if we are not to treat people as equal in dignity and not mere objects of others' ends. Ultimately what we protect, how we protect it and how different goods and values are balanced should be in accordance with people's informed preferences, not as a matter of individualised notice and consent but as a matter of collective self-government.

We therefore need to draw on other mechanisms for public engagement outside the electoral cycle to ensure that legitimate concerns about data practices are heard and addressed.

This is not an original observation. In 2012, the Commonwealth government published a document entitled *Science and Technology Engagement Pathways: Community involvement in science and technology decision making*. This document, no longer on any government website, commenced with the observation that:

In democracies, there is a recognition that citizens should have input into decisions that affect them.

Communities are consulted about city planning, regional development and infrastructure projects like roads and waste facilities, so why not new developments in science and technology, which may affect them just as much?

The document itself prescribed a public engagement framework involving seven principles: commitment and integrity, clarity of objectives and scope, inclusiveness, good process, quality information/knowledge sharing, dialogue and open discussion, and impact on decision-making.

The technology hot topic at the time was nanotechnology rather than AI, but the principles themselves are context independent. This document is of course only one tiny piece of a broad international literature on public engagement, including in relation to technology assessment, but having been devised by a multistakeholder process in Australia, it is a useful place to start.

We could set up a series of opportunities for Australians to learn about current public sector and private sector data practices and share their concerns, taking inspiration, perhaps, from some of the activities of the Centre for Data Ethics and Innovation in the UK. Because there is no single right answer about how opportunities and harms associated with modern data practices might be balanced, discursive online and local fora would provide an opportunity to gauge priorities.

Conversations will need to be inclusive, particularly with respect to vulnerable populations who often experience harms associated with automated data processing. While lawyers such as ourselves can outline minimum requirements (such as the rule of law, human rights, and compliance with existing legislation), in a democracy it is important that the discussion of "where to from here" is a broad one.

In short, our argument here is that conversations about data and data dilemmas have so far been too narrow, both in terms of who is involved, and in the framing around privacy and 'privacy versus' other goods.

We have far too often ended up talking about data use as a privacy issue, and then been reduced 'balancing' privacy against security, or health, or innovation, as if data policy were a seesaw, where one of the two must be ascendant.

We conflate people's preferences and concerns about data collection and use with people's concern for privacy, which turns data policy into an all or nothing: either people want privacy (and aren't on Facebook) or not (and then they're on Facebook, and it's a free for all).

Broadening the conversation about data practices both to bring in a wider range of people and perspectives, and also to think more expansively about what might be at stake for people, could be a way to help government and corporate actors understand that what people might want protection from is not just collection or revelation of information about them.

People's concerns about data and data use are also about autonomy and manipulation; or being treated as the object of someone else's activities; about arbitrary differential treatment and the potential abuse of power. Finding ways to reflect this more nuanced set of concerns around data and its use is going to require conversation and dialogue, and not just one, but ongoing conversations in a wide range of contexts.

Systems

The third level at which we need to be building in respect for the rule of law and human dignity is at the level of systems. By 'systems' we mean networks of humans and technologies that perform particular tasks. For example, the 'robodebt' system comprised data matching tools, software that performed calculations (that did not always get the right answer), printers that produced letters, web platforms that were difficult for those affected to access and use, and humans who had very little control over system outputs.

Designers of systems that collect and use data to undertake functions or make decisions that affect people necessarily make a series of choices: about what and how data will be collected; what technology will be used to analyse it; and what the outputs of the system will look like, how they will be communicated to people, how much human involvement and control there will be, and what systems will be used for correction or contestation.

Systems can be designed to be more consistent with the rule of law. For example, a core demand of the rule of law is that decisions with a significant impact on people's rights can be justified to the people affected. This is one way we ensure that the exercise of power is not (and is not seen to be) arbitrary.

If we are going to incorporate data-driven predictive analytics into our decision-making, for example, a determination as to how an offender serves their sentence, we should not be relying on techniques (such as random forests and deep neural networks) that cannot be adequately explained. And accountability will require human involvement in many kinds of decisions.

Similarly, system design can be conducted with human dignity in mind. Privacy is an aspect of this: we might actively decide to adopt, for example, privacy enhancing technologies, understanding that protecting information from unwanted revelation is an important part of respecting human dignity.

But it is only a part of the picture, as we have sought to emphasise, and more broadly, we might want to demand that consumer-facing or citizen-facing systems show some minimum level of respect for human autonomy. This may be by, for example, avoiding some forms of manipulation that undermine conscious decision-making, whether it is designed to make people buy things (Manwaring 2018) or to shape behaviour in ways that the government has decided is 'for people's own good'.

As for *how* we bring these principles into system design, we would suggest that this is where a risk-based framework could be of assistance. There are plenty of people talking about risk assessment frameworks in the context of data problems. Privacy impact assessments are, in essence, a risk assessment framework focused on existing privacy legislation.

Scholars like Metcalf et al. (2021) and the AI Now Institute via its proposal for algorithmic impact assessments have expanded on this method, suggesting an assessment according to the impact of these technologies on (some combination of) individual or collective risks to human health, property, the environment or fundamental rights.

The AHRC in its *Human Rights and Technology* final report (2021) recommended that the government should legislate to require a human rights impact assessment before government uses AI-informed decision-making for administrative decisions, and that government should encourage similar processes by private sector entities (recommendations 2 and 9). The EU's proposed AI Act is significantly built on requirements for conformity and risk assessments for entities proposing to use AI systems designed as 'high risk'.

But, as argued above, the technologically defined boundaries of such laws are often arbitrary. Governments and corporations should embed human rights thinking into core policies and functions, not hand the issue over to the IT department. There may be specific issues in the context of AI, and technical people will need to be involved in ensuring systems are designed with human rights in mind. Some technologies may simply not pass the test, either generally or in specific contexts (Pasquale 2018). The point is, however, that this should be an implication of a broader law and policy goal, rather than something crafted independently.

The strength of these proposals is that they move beyond current privacy impact assessments that are built on traditional privacy law. We would add two nuances to these proposals. First, we would argue that any risk assessment framework ought to be based not on the kind of technology used (whether AI is used or not, or whether decisions are based 'solely' on automated decision-making) but on the degree and nature of the impact of the system (human and technical) on people, society and the environment.

In our view, the European approach is flawed because its scope is limited by a technological frame. While it distinguishes between systems that are and are not 'high risk', it allows equally harmful non-AI systems to flourish. The problem for both the European legislation and some of the recommendations by the AHRC is that they start with problems defined around a set of technologies rather than with the broader frame. Their recommendations are thus useful but also incomplete.

Our second nuance is that the risks considered should include examination of both the rule of law and human dignity and autonomy. Confining any impact assessment, as the AHRC does, to a human rights impact assessment arguably leaves out the

important considerations we have discussed, and leads to a temptation to focus on known human rights risks (most likely discrimination and privacy risks). Our suggestion also goes beyond the EU proposal's risk assessment, which focuses on risks of harm to the health and safety or adverse impact on the fundamental rights.

Conclusion

Ultimately, the questions we have been looking at in this chapter require political and societal resolution, not just a dialogue of experts. Our purpose here is to begin a conversation, with our fellow Australians in particular, about what kind of legal regime, and what kind of principles and considerations built into our legal regimes, might offer the best protection from harms resulting from new and emerging data practices.

After considering the limits of existing legal frameworks, even in an amended form, we suggest some other possible starting places – different kinds of legal regimes with different purposes.

We focused on the rule of law, and considerations of human dignity and autonomy, because in our view, while also incomplete, they help us to understand why it is troubling for public or private sector to put people in categories based on their characteristics, and then, on the basis of those characteristics, shape what they see, what they can know, and what they are offered.

But in the end, in order to understand the sufficiency of any of these approaches, alone or in combination, we need to better understand the kinds of concerns that people have and the relative importance of different things. In other words, we need to broaden public understanding and ignite public debate about why and how data practices might be channelled in productive and protective ways.

References

- Angwin J, Larson J, Mattu S and Kirchner L (23 May 2016) 'Machine Bias', ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Australian Competition and Consumer Commission (2019) Digital Platforms Inquiry: Final Report. <https://www.accc.gov.au/publications/digital-platforms-inquiry-final-report>.
- Australian Human Rights Commission (2019) Human Rights and Technology: Discussion Paper. <https://humanrights.gov.au/our-work/rights-and-freedoms/publications/human-rights-and-technology-discussion-paper-2019>
- Australian Law Reform Commission (June 2014) Serious Invasions of Privacy in the Digital Era (ALRC Report 123). <https://www.alrc.gov.au/publication/serious-invasions-of-privacy-in-the-digital-era-alrc-report-123/>
- Barocas S and Nissenbaum H (2014) 'Big Data's End Run around Anonymity and Consent', in Lane J, Stodden V, Bender S and Nissenbaum H (eds) *Privacy, Big Data and the Public Good*, Cambridge University Press.

- Bennett Moses L (2020) 'Who Owns Information? Law Enforcement Information Sharing as a Case Study in Conceptual Confusion', *University of New South Wales Law Journal*, 42(2):615–41.
- Cannataci J (1987) *Privacy and Data Protection Law*. Norwegian University Press.
- Centre for Data Ethics and Innovation (4 February 2020) *Online Targeting: Final Report and Recommendations*. <https://www.gov.uk/government/publications/cdei-review-of-online-targeting/online-targeting-final-report-and-recommendations>
- Cohen JE (2019) *Between Truth and Power: The Legal Constructions of Informational Capitalism*, Oxford University Press. <https://doi.org/10.1093/oso/9780190246693.001.0001>
- Dencik L, Redden J, Hintz A and Warne H (2019) 'The "Golden View": Data-Driven Governance in the Scoring Society', *Internet Policy Review*, 8(2):1–24. <https://doi.org/10.14763/2019.2.1413>
- Department of Industry, Innovation and Science (2019) 'AI Ethics Principles', Australian Government. <https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles>
- Department of the Prime Minister and Cabinet (2019) *Data Sharing and Release Legislative Reforms Discussion Paper*. https://apo.org.au/sites/default/files/resource-files/2019-09/apo-nid259016_1.pdf
- Hanrahan C (16 April 2019) 'Environment Trumps Economy as Voters' Top Concern Ahead of Election'. <https://www.abc.net.au/news/2019-04-17/vote-compass-election-most-important-issues/11003192>
- Fjeld J, Achten N, Hilligoss H, Nagy A and Srikumar M (2020) 'Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-based Approaches to Principles for AI', Berkman Klein Center for Internet & Society. <https://doi.org/10.2139/ssrn.3518482>
- Gaze B and Smith B (2017) *Equality and Discrimination Law in Australia: An Introduction*, Cambridge University Press.
- Goggin G, Vromen A, Weatherall K, Martin F, Webb A, Sunman L and Bailo F (2017) *Digital Rights in Australia*. <https://ses.library.usyd.edu.au/handle/2123/17587>
- Greenleaf G (2019) 'Global Tables of Data Privacy Laws and Bills', Supplement to 157 *Privacy Laws & Business International Report*.
- Johnston A (2020) 'Individuation: Re-Imagining Data Privacy Laws to Protect against Digital Harms', *Brussels Privacy Hub Working Paper Series*, 6(24):22.
- Krygier M (2011) 'Four Puzzles about the Rule of Law: Why, What, Where? And Who Cares?', *Nomos*, 50:64–104.
- Krygier M (2019) 'What's the Point of the Rule of Law', *Buffalo Law Review*, 67:743. <https://digitalcommons.law.buffalo.edu/buffalolawreview/vol67/iss3/16>
- Leonard P (2020) 'Data Privacy in a Data- and Algorithm-Enabled World', *Privacy Law Bulletin*, 17(3):43–47.
- Luguri J and Strahilevitz L (2019) 'Shining a Light on Dark Patterns', University of Chicago Coase-Sandor Institute for Law & Economics Research Paper 879, Social Science Research Network. <https://doi.org/10.2139/ssrn.3431205>
- Manwaring K (2018) 'Will Emerging Information Technologies Outpace Consumer Protection Law? – The Case of Digital Consumer Manipulation', *Competition and Consumer Law Journal*, 26(2):141.

- McCrudden C (2008) 'Human Dignity and Judicial Interpretation of Human Rights', *European Journal of International Law*, 19(4):655–724. <https://doi.org/10.1093/ejil/chn043>
- Mendoza I and Bygrave LA (2017) 'The Right Not to Be Subject to Automated Decisions Based on Profiling', in Synodinou TE, Jougoux P, Markou C and Prastitou T (eds), *EU Internet Law: Regulation and Enforcement*, Springer International Publishing. https://doi.org/10.1007/978-3-319-64955-9_4
- Metcalf J, Moss E, Watkins EA, Singh R and Elish MC (2021) 'Algorithmic Impact Assessments and Accountability: The Co-construction of Impacts', in FAcCT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 735–746. Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445935>
- Nissenbaum H (2004) 'Privacy as Contextual Integrity', *Washington Law Review*, 79(1):119–158.
- Obermeyer Z, Powers B, Vogeli C and Mullainathan S (2019) 'Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations', *Science*, 366(6464):447–453.
- Ohm P (2010) 'Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization', *UCLA Law Review*, 57:1701–1777.
- Pardes A (2 August 2020) 'How Facebook and Other Sites Manipulate Your Privacy Choices', *Wired*. <https://www.wired.com/story/facebook-social-media-privacy-dark-patterns/>
- Pasquale F (18 June 2018) 'Data Nationalization in the Shadow of Social Credit Systems', LPE Project. <https://lpeproject.org/blog/data-nationalization-in-the-shadow-of-social-credit-systems/>
- Richards DAJ (1981) 'Rights and Autonomy', *Ethics*, 92(1):3–20.
- Sex Discrimination Act 1984 (Cth). <https://www.legislation.gov.au/Details/C2014C00002>
- Trakman L, Walters R and Zeller B (2019) 'Tort and Data Protection: Are There Any Lessons to Be Learnt?', *EDPR Review*, 5(4):1–20.
- Veale M and Edwards L (2018) 'Clarity, Surprises, and Further Questions in the Article 29 Working Party Draft Guidance on Automated Decision-Making and Profiling', *Computer Law and Security Review*, 34(2):398–404. <https://doi.org/10.1016/j.clsr.2017.12.002>
- Viljoen S (11 November 2020) A Relational Theory of Data Governance (SSRN Scholarly Paper No ID 3727562, Social Science Research Network). <https://papers.ssrn.com/abstract=3727562>
- Whitman JQ (2004) 'The Two Western Cultures of Privacy: Dignity versus Liberty', *Yale Law Journal*, 113(6):1151–1221. <https://doi.org/10.2307/4135723>.
- Zalnieriute M, Bennett Moses L and Williams G (2019) 'The Rule of Law and Automation of Government Decision-Making', *The Modern Law Review*, 82(3):425–455. <https://doi.org/10.1111/1468-2230.12412>

Chapter 4

Trust building for data sharing – Understanding trust as a social relationship



By Theresa Anderson

Theresa Dirndorfer Anderson is the Director and Social Informaticist at Connecting Stones. A social informaticist with a PhD in Information Science, her award-winning work as an educator and researcher engages with the ever-evolving relationship between people and emerging technologies. Theresa now focuses on advancing socially-just data policies and building trusted environments for data/AI use. She contributes to the development of reference and actionable frameworks at local and international levels. In 2021, she was appointed to the NSW Government's inaugural Artificial Intelligence Advisory and Review Committee. She is actively contributing to development of an international standard for Data Usage (ISO JTC1/SC32/WG6), serving as a Project Editor. Theresa also regularly contributes to International Science Council (ISC) Committee on Data (CODATA) initiatives enhancing global cooperation on FAIR data policy and practice.

Careful wintertime

A little, green frog jumps on

Trusting in the stone

Trust is an important component of our system of governance. It is considered a keystone for building and maintaining a flourishing modern urbanised society (Beck 1992, 2000; Luhmann 2018). Global communications firm Edelman has been studying trust for 20 years, sharing its findings through the Edelman Trust Barometer and ongoing global surveys. Its 2020 and 2021 findings on trust in business, government, media and NGOs reveal an erosion of trust in all four sectors, which Edelman attributes to:

people's fears about the future and their role in it, which are a wake-up call for our institutions to embrace a new way of effectively building trust: balancing competence with ethical behaviour. (<https://www.edelman.com/trust/trust-barometer>)

In the ecosystem of data sharing and use, building and maintaining public trust is essential for maintaining public confidence in the way that data (especially data taken from the public) is being used. As the previous chapters explain, we also see rising concerns about the governance around data sharing. This is needed so data can be channelled productively without diminishing the essential protections the public is entitled to expect.

Building on these earlier discussions, this chapter offers:

- theoretical framings to use as markers to help with sensemaking about trust relations
- lessons for the future via a brief sociohistorical snapshot of the crisis of trust
- a framework for understanding keystones practices for building and maintaining trust relations
- opportunities for engagement to co-design trust-building frameworks.

Understanding the conditions that help us to build and to demonstrate trust in our meaning-making and in our social relations lays the foundations for designing trustworthy data and AI-enabled technologies.

Understanding trust

Trust is a way to control everyday interaction with the future (Luhmann 2018). While trust is not the sole foundation of the world, Luhmann suggests the world as we know it could not function without it:

Trust, in the broadest sense of confidence in one's expectations, is a basic fact of social life. In many situations, of course, a person can choose in certain respects whether or not to bestow trust. But a complete absence of trust would prevent him or her from even getting up in the morning. He would be prey to a vague sense of dread, to paralyzing fears. He would not even be capable of formulating definite distrust and making that a basis for precautionary measures, since this would presuppose that he trusts in other ways. Anything and everything would be possible. Such abrupt confrontation with the complexity of the world at its most extreme is beyond human endurance.

Trust is indispensable in a social system, but because it is highly situational and contextual it is never values-neutral.

What does it mean to trust?

While it is beyond the scope of this chapter to cover all the many definitions and perspectives on trust, it is important to recognise that in both theory and practice, trust is described as an elemental feature of our social worlds. Jaffe (2018) for instance refers to it as 'the glue of society':

Its presence cements relationships by allowing people to live and work together, feel safe and belong to a group. Trust in a leader allows organizations and communities to flourish, while the absence of trust can cause fragmentation, conflict and even war. That's why we need to trust our leaders, our family members, our friends and our co-workers, albeit in different ways.

Both as actions we perform and as an object we value, trust is about belief, confidence, reliability and a sense of truth (see, for instance, Oxford English Dictionary definitions).²⁴ Trust is a bond that links us to other people; it is a sentiment that:

lets us put greater confidence in other people's promises that they mean what they say when they promise to cooperate. (Uslaner 2003:43)

Trust is a leap into the future. The predictive nature of trust means we must recognise it not as a stable concept across time but rather as one that is:

fluidly interpreted across locations and used as a focus for elaborating concerns, complete lack of concern and the identification of who should be concerned in relation to new systems. (Neyland 2006: 151)

While an organisation's (or a government's) reputation reveals insight about past performance, trust is more forward-looking.

²⁴ Extracts from Oxford English Dictionary Online:

TRUST (n) a. Firm belief in the reliability, truth, or ability of someone or something; confidence or faith in a person or thing, or in an attribute of a person or thing. Chiefly with *in* (formerly also *†of*, *†on*, *†upon*, *†to*, *†unto*). b. The quality or condition of being trustworthy; loyalty; reliability; trustworthiness.

TRUST (v) To have faith or confidence in a person, quality, or thing; to rely on.

However, even as a future-oriented construct, a judgement about trust is shaped as much by what you do, or have done, as it is about how you do it. The Edelman Trust Barometer frames trust as being granted on two distinct attributes:

1. competence (delivering on your promises and how well you get things done)
2. ethical behaviour ('doing the right thing', working to improve society, honesty and fairness).

To assess the 'trust capital' of an organisation, Edelman's framework uses a metric called the Edelman Net Trust Score (ENTS) calculated across four key dimensions that drive corporate trust: ability, integrity, dependability and purpose (where purpose is related to effort to have a positive impact on society).

An organisation's ENTS score is based on the analysis of stakeholders' responses to the question: to what extent do you trust the organisation to do what is right? (R Edelman 2020).

The Organisation for Economic Cooperation and Development (OECD 2022) points to similar drivers of trust in government institutions, using its own survey of trust²⁵ to capture the degree to which institutions are responsive and reliable in delivering policies and services, and the degree to which act in line with the values of openness, integrity and fairness.

Trust is a decision-making process, where a judgement links past actions (and reputation) to future potential actions. Given that no future can be certain, such judgements carry an element of risk. The more trusted the relationship between an organisation and its stakeholders, for instance, the more risk is likely to be accepted. Conversely, in a low trust environment, taking leaps into an unknown future brings a greater sense of risk.

In this way, trust is linked to risk-taking. It is what allows us to move beyond doubt and into a more productive and positive engagement with the unknowns of our worlds – in the present and in our possible futures. We are more likely to tolerate the uncertainty of any situation we face when we have a sense of trust about the people or setting involved. In the early stages of the declaration of the global COVID-19 pandemic, consulting firm McKinsey made similar observations:

In crises, the state plays an essential and expanded role, protecting people and organizing the response. This power shift transforms long-held expectations about the roles of individuals and institutions. (Craven et al. 2020)

In light of this erosion of trust, Craven et al. argue, a rethinking of the social contract is taking place. These comments are an impetus for this chapter's driving question focusing on trust as a social value: how do we generate sufficient trust to allow us to move forward as a community, given that the landscape will continue to change as we seek to put protections in place?

25 <https://www.oecd.org/governance/trust-in-government/>

Why trust matters: Trust in uncertain times

The growing crisis of trust discussed in the media and experienced in our daily practices is showing how critical trustworthy leadership and trusted information channels can be in multichannel, mediated environments:

When the velocity of progress increased beyond a certain point, it becomes indistinguishable from crisis (Barfield 1993:152).

Eisenberg (2001:550) observes we need to 'develop new ways of living in a world without foundations'. The challenge comes in terms of overcoming the fear of the future and all the unknowns in a time when there is so much flux.

To understand why trust matters -- particularly in times of crisis -- we must first understand how it is connected to risk, uncertainty and the ability to move forward. At both the individual and the system level, 'trust depends on the inclination towards risk being kept under control and on the quota of disappointments not becoming too large' (Luhmann 2018:98).

Wallerstein's (1998) observations about global complexity and human social systems in periods of transition suggest that fear and panic kick in when we perceive our situation to be precarious, individually and collectively.

Wallerstein argued that fear can be brought on by the major impact that seemingly small inputs can have on our stability, leading to a sense of crisis. The cascading effect of the subprime credit crisis earlier this century, job losses accompanying digitisation and automation strategies, and our current global trust deficit all seem to confirm his assessment.

In the current post-pandemic climate, trust, risk and uncertainty permeate media reports and even everyday conversations; the pursuit of 'certainty', 'assurance', 'reasonable risk' and 'measured risk' seem to be increasingly sought after in so many parts of our lives and our society. Beck (1992) refers to risk as the main feature of modern society. And yet, as Åsa Boholm observes:

In real life situations, the boundary between certitude and uncertainty is of course seldom razor-sharp, and vagueness and ambiguity tend to be the rule rather than the exception. (Boholm 2003:168)

Uncertainty and risk are often grouped together in our conversations and imaginings. How we frame uncertainty and risk has much to do with where we are at a given moment and what we are experiencing. Eisenberg (2001:534) goes so far as to suggest that our primary challenge as human beings is 'living in the present with the awareness of an uncertain future'. Anderson (2006) has shown that positive and negative forms (as experienced at any one moment) are inextricably intertwined, but one key to working through any kind of uncertainty is developing a tolerance for it.

Risk and uncertainty are linked – when something of value is at stake, uncertainty

can relate to the chances of a negative outcome and the nature of the outcome itself. People work on ways to overcome, manage or deal with the uncertainty and risk experienced in their lives on a daily basis. Malaby positions risk and uncertainty as an important element of our sociality:

It is through the engagement of indeterminacies, rather than their minimization or resolution ... that one may socially demonstrate one's place vis-à-vis chance, and by extension, one's place in relation to others in the world. (Malaby 2002:284 as cited by Christensen and Mikkelsen 2008:113)

This perspective is a critical departure from assuming risk is necessarily dangerous or destructive, or that uncertainty needs to be avoided or eliminated. This anthropological stance helps us to appreciate that everyday life is characterised by uncertainty – and that uncertainty can bring unexpected pleasures as well as pain.

Critically for this discussion, however, we must appreciate that it is the perception of risk or uncertainty that is at issue. Both are socially constructed phenomena, intersubjectively produced and culturally located.

Beck (1992) describes risk as a state between security and destruction. He is not suggesting that risk has to be negative, but many discussions drawing on his construct of the risk society tend to suggest that it is a pessimistic view, because risk is a sign of trouble and trauma. Beck takes issue with the pessimistic interpretation people have made of his conception of the risk society, pointing to opportunities of the 'bads' (Beck 2000:226). Even in the bad there is opportunity, and it is that opportunity that is worth examining in relation to human responses to uncertainty.

Boholm (2003) and Malaby (2002) further push this argument that we need to view risk in more nuanced ways, as neither simply objective nor subjective. Boholm, for instance, draws on a sociological definition of risk as:

a situation or event where something of human value (including humans themselves) has been put at stake and where the outcome is uncertain. (Rosa 1998:28 as cited by Boholm 2003:166)

If we adhere to Boholm's claim that uncertainty is a necessary feature of our existence, then we need to acknowledge that working with and through uncertainty is an everyday experience. More than that, there is much evidence to suggest that uncertainty and curiosity are so closely linked that were we to reduce the uncertainties in our worlds, we would in effect be closing the door on the opportunities for innovation and creativity that are desired in so many sectors of our society.

It is important to recognise that there are risks and uncertain situations that seem to have little creative potential. Risks to family security (for example, job loss, housing concerns), to health (for example, surgery, illness), or to personal security (for example, crime, terrorism) are examples that many of us can appreciate. However, even in such circumstances, individual judgements vary as to where to draw the line in terms of threats to our security and acceptable risks.

A review of research into terrorism threats, for instance, found great variance in terms of the perception of risk and potential terrorist threats within different communities at different points in time (Maguen, Papa and Litz 2008).

Risk reaps reward and so there are many examples of the productivity of risk in human practice. There are people who make the choice to actively engage with risk, some physical (for example, extreme athletes) and some economic (for example, financial speculators).

Zaloom (2004) explored the productive life of risk through fieldwork on the trading floor of the Chicago Board of Trade, a major global financial futures exchange. Zaloom describes the 'fine balance necessary to work with risk' and explains that it involves working with norms of risk management that are generated (in her context) on the trading room floor to the extent that self-definition and group formation coevolve: 'Active engagements with risk are a locus of self and space in contemporary economic and social life'.

Here we find a conundrum of risk and uncertainty and the mixed bag of perceptions of both across situations and cultures. The concept of risk can be understood as a framing device – allowing us to transform it from 'an open-ended field of unpredicted possibilities into a bounded set of possible consequences' (Boholm 2003:167).

Risk can be conceptualised and managed in different ways across communities, cultures and organisations. Looking at the productivity of risk draws attention to ways that some people see it fitting into their work and their self-defining behaviours.

Perception is a powerful determinant when it comes to developing a tolerance of risk and uncertainty in society collectively and in our own lives. There is a powerful social element at play in the way we approach risk, for instance. Our sense of self and the way we wish to see ourselves in relation to particular social groups informs the way we approach risk and uncertainty.

Thus, we can see uncertainty as a fundamental experiential realm of human existence associated with tolerance and risk-taking. It is through the experience of risk and uncertainty that we learn to identify how much we can individually endure.

It is how we develop resilient capacities to tackle future challenges. While uncertainty or risk are not inherent, research suggests that when an individual threshold is reached, the negative emotions can overpower us. At moments when uncertainty and risk seem too much to bear, these powerful emotions diminish any opportunity for exhilaration and the pleasures of uncertainty (Wilson et al. 2005). This anthropological perspective helps us to appreciate that everyday life is characterised by uncertainty – and that uncertainty can bring unexpected pleasures as well as pain.

Taking a holistic view of uncertainty, we begin to appreciate that thresholds exist along a certainty-uncertainty spectrum. Too much 'not knowing' can overwhelm and lead to the frustration associated with information overload, unmanageable uncertainty,

and risks beyond the tolerable. The uncertainty that Beck (1992, 2000) appears to be talking about, for instance, involves a lack of knowledge and challenges in terms of how information will become available and whether it can become available.

Understanding this interplay between risk and uncertainty and our perception of the risk landscapes we experience is essential for appreciating why trust matters so much in times of (perceived) crisis. Helping to maintain this delicate balance of uncertainties and pain points (metaphoric and real) – at both individual and collective levels – is the level of trust we have in ourselves, in our families, in our leaders, in data. Trust addresses fear and supports moving forward into the unknown.



The image presented here uses the idea of windsurfing as a means of illustrating the significant role that perceptions of both uncertainty and vulnerability play in our judgements of trust. A personal inclination towards the adrenaline rush of such an adventure sport will determine how much uncertainty a windsurfer might be prepared to accept in any one incident on the waves. This confidence is in part determined by the trust they have in their own expertise as a windsurfer, trust in their knowledge of the conditions that day, and trust in their equipment.

Equally, the more vulnerable one feels in such a situation, the more evidence of trustworthiness in the factors at play in that situation one will likely need. Whatever challenge we might face, moving forward sees this interplay between our judgements of uncertainty, vulnerability and trust.

In summary, trust underpins social order. It is part of the process that mobilises both rational and emotional components of human judgement and value systems, drawing on both direct and indirect experiences. It is also a means for building equity and justice into the social fabric (Uslaner and Brown 2005). Trust is not easily earned but can be very quickly eroded. Critically, as decision-makers and designers of data technologies, it is essential that we appreciate that trustworthiness is not a single step but an ongoing cycle of hard work to gain, maintain and demonstrate being worthy of a community's trust.

How and why we learn to trust

Trust may be difficult to define with precision, but we do know when it's lost. When that happens, we withdraw our energy and level of engagement. The erosion of trust brings its absence front and centre into conversation and into our lives. Distrust is closely linked to discord.

Conversely, in a context where trust is considered to be 'high', we are likely to pay less attention to the presence of this trust relation, because we are actively engaged in the other things that are happening. Like the windsurfer depicted in the image above, we would be actively engaged in the pursuit and enjoying the moment. This, then, is how we help build a flourishing community.

We become a 'trusted person' (or entity) through the actions we take. To some extent, the reputation we've built as a trusted person can signal to others that we are worthy of trust. However, reputation is an outcome of the past, whereas trust implies a projection. We understand and enact 'trust' through breaches, expectations and repairs of (social) world/order:

Trust underpins or makes possible social action and makes possible expectations of consistent future social action rendering a social order possible and meaningful. (Neyland 2006:161)

A child, for instance, will test the boundaries of how far they can go until they have been scolded for going too far. The dynamism that is inherent in this relationship is important, flagging again that there are few absolutes. When there is a breach and recognition one has gone too far, steps are taken to repair the situation.

Negotiating trust and rebuilding it after a breach are not new social practices, as the etymology of the word 'trust' demonstrates. However, as our interactions became increasingly mediated by technology, our social relations increasingly spread across time and space to become 'disembedded' from their local contexts (Giddens 1990). And while disinformation and misinformation are also not novel, our contemporary communication tools contribute to the rise of 'mythinformation' (Winner 1984; Burke 2020), media mythmaking (Huff and Rea 2009) and infodemics (Minors 2021).

A state of mistrust can erupt, especially when the legitimacy of decision-makers and the process of decision-making is questioned – as has been the case throughout two current crises of the global pandemic and climate change:

It is only through engaging deliberately with each other and with the facts that we can learn to trust and share the information that keeps us alive. (Minors 2021:27)

Furthermore, as Taddeo (2009) discusses, the emergence of trust in digital contexts has created new theoretical problems. Consequently, in the many mediated engagements we experience in our increasingly digitised and datafied existence, ascertaining the trustworthiness of a claim made by a person or an organisation becomes a more deliberate design requirement, involving more deliberate engagement with the community.

Earning trust

When we are discussing trust relations, we are really looking at judgements of trust and trustworthiness, either as someone looking for evidence of the trustworthiness of another person (or entity) or seeking to reassure someone else that we are ourselves worthy of their trust.

Building on this understanding of trust and trustworthiness, I suggest that – as individuals and as organisations – we build trust in four quadrants, depicted in the following figure:



Reassurance:

Communication and professionalism



Resilience:

Persistence and creativity



Relationships:

Building and maintaining connections



Reflection:

Time to think and test

The following sections describe each of these four quadrants further.

Reassurance: Communication and professionalism

Building on our earlier description of trust as a forward-looking concept, we can think about trust as a powerful KPI that is a forward-looking projection of our intentions for the future. The reassurance that stems from the competence and professionalism of medical professionals and scientists working on COVID-19 vaccines and in hospitals, for instance, has been frequently invoked during the global pandemic to help build trust in the actions taken by authorities.²⁶

For KPMG (2018), for example, trust in data and analytics is founded on four key anchors:

1. quality of the data, models and algorithms
2. effectiveness – extent to which analytics deliver the desired results
3. integrity – ethical and acceptable use
4. resilience – optimising data and analytics for the long term.

As this 2018 report goes on to explain, these anchors are then used to communicate the professionalism of work undertaken.

Reassuring a community about the professionalism applied to any situation, however, also brings with it a responsibility to contribute to the literacies of that community so that they can be sufficiently informed and active in public evaluations of such practice. A commitment to public education is very important for this quadrant, especially in democratic society.

²⁶ See, for instance, Jonathan Watts's June 2020 interview with Bruno Latour about expertise in times of crisis (<https://www.theguardian.com/world/2020/jun/06/bruno-latour-coronavirus-gaia-hypothesis-climate-crisis>).

Relationships: Building and maintaining connections

Trust is about relationships: the relationship between our words and our actions; between our past practices and our stated future actions; and critically, between the two parties negotiating a trust relation.

When we talk about whether or not we find an individual trustworthy, we are likely making that judgement based on our relationship with that person and what is known (or unknown) about them. Even at the organisational level, we will still be thinking about relationships with that entity or individuals representing that entity. So, at its core, when someone makes a judgement about trustworthiness in any given situation, they are still ultimately drawing on local experience. This is precisely why trust and risk go hand in hand – the more trust there is at a given point in time in a relationship (with an individual or an organisation), the more risk is going to be accepted.

Results from Edelman Trust Barometer surveys reveal how critical local connections are for people across the globe. Employees expect to be heard in their organisations, with 73% of all respondents in the 2020 survey results expecting any prospective employer to not only give them an opportunity to have a say in shaping the future of society but also include them in the organisation's planning for the future.

The Edelman Trust Barometer findings also identified more trust in local/state government than in central/federal government. In 18 out of the 24 international markets analysed as part of their 2020 report, for instance, local government was more trusted (R Edelman 2020:41). Going back to the earlier observation about vulnerability and perceptions of uncertainty, perhaps this finding has something to do with the fact that as humans, we feel particularly vulnerable at moments of crisis and thus value contacts closer to home.

Resilience: Persistence and creativity

Resilience is about being able to handle challenges and manage risks so we can learn as we go – even from our mistakes. Tenacity, persistence and experimentation are very important qualities when faced with challenges and seeking to find a way to move forward. We can also begin to see the linkages between judgements of trust, risk and resilience (for example, Anderson 2006, 2013, 2020). Building trust is critical to create a resilient society amid imperfect and incomplete information.

Reflection: Time to think and test

There is also great power in the pauses we can introduce into our practice (Anderson 2013). Learning from past actions (both successful and not) as well as ensuring reflection on present activities can make it possible to plan better in future. Learning, for example, from past breaches as well as successes and then making that learning visible to the community is a powerful tool for demonstrating trustworthiness.

The next section uses the story of Ignaz Semmelweis (seen by many as the father of infection control and hand hygiene) as a way to operationalise these four quadrants

in which we can build trust (resilience, reassurance, relationships and reflection) and to illustrate their interplay in action.

Trust building is personal, local and political

As Neyland (2006) reminds us, trust is not a universal or stable concept that can be rigidly hooked to a single definition to apply across multiple contexts. To appreciate the important role that relationship building can have and to help us understand our engagements with emerging technologies in times of crisis, let's examine these four quadrants of trust building via the story of Dr Ignaz Semmelweis in 19th Century Vienna.

Accounts of how he sought to resolve a problem in the Viennese hospital where he was working, and the method that he used, illustrate how expertise and authority are insufficient for introducing new ideas without putting effort into local knowledge and relationships.

Semmelweis approached the challenge of unexplained infections in his hospital by testing out and systematically ruling out what was thought to account for uncontrollable deaths in one wing of the hospital. It was only through trial and error that he ultimately identified that handwashing as part of higher-level infection control could make the difference.

Early on, he theorised that the surgeons and the interns who going from autopsies into the hospital to help with the delivery of babies were associated with the unexplained deaths. Initial theorising posited that these specialists were carrying disease particles from one place to the next. However, ultimately through trial and error of different treatment methods, Semmelweis worked out that using a new approach to handwashing could resolve the problem. His innovative method did dramatically reduce the deaths down to less than 1%.

This idea of *recognise-explain-act* is an approach that has remained in place in public health throughout the current global pandemic (World Health Organisation, 2009). Put simply, it means that when you encounter a challenge or a difficulty, you must first try to recognise what the problem is, then you test the waters through trial-and-error explorations seeking to explain causations, and then take action.

Unfortunately, Semmelweis had great trouble convincing the people he worked with about the benefits of continuing with his procedure. Some of them did it in the early days, but because it was a new practice that was onerous and tedious (in the eyes of the staff asked to perform this new task), it was difficult to consistently and sustainably implement the change.

Despite this new knowledge he'd uncovered using his professional training as a doctor and scientist, he did not have good working relationships with the people around him. Because there was so little respect for him in that local context, his approach was not readily adopted by his peers:

Semmelweis experienced great difficulties in convincing his colleagues and administrators of the benefits of this procedure. In the light of the principles of social marketing today, his major error was that he imposed a system change (the use of the chlorinated lime solution) without consulting the opinion of his collaborators. Despite these drawbacks, many lessons have been learnt from the Semmelweis intervention; the “recognize-explain-act” approach has driven many investigators and practitioners since then and has also been replicated in different fields and settings. (World Health Organisation 2009: Chapter 4)

This account of innovation in times of crisis is very similar to accounts of the cholera epidemic in London and the challenge that Reverend John Snow faced, where a personal connection was also lacking thereby blocking acceptance and trust in his approach, no matter how much scientific evidence he could muster (Tulchinsky 2018:77-99).

Even though scientifically Semmelweis demonstrated that what he was doing was appropriate and delivered results (saving lives), he struggled to convince others of his expertise and the validity of this new intervention he had designed. He didn't have relationships with the other colleagues in his hospital, which blocked him from putting his new initiatives into practice.

Connections and communication go hand in hand with professionalism to build a trusted relationship. This account of Semmelweis's struggle to convince his peers to implement the new approach he'd developed through rigorous investigation shows that we must value relationship building and maintaining those connections at that local level.

It illustrates not only the importance of relationships for becoming a trusted person, but also the necessity for resilience. Semmelweis had to be persistent in his method of inquiry and in his efforts to convince others that he had a trustworthy approach to offer them. It was through trial and error, curiosity, persistence and tenacity that change occurred.

Building trust through a culture of care

Recognising the personal, perceptual character of trust helps us appreciate how notions of care, compassion, empathy and wellbeing can act as drivers for building trust. Anderson (2020:186-187) explores the link between vulnerability, uncertainty and trust in greater detail, framing a set of keystone practices for building a trusted environment for data use:

- community: reflecting an appreciation for the interdependencies and complexities of all constituent parts
- civility: showing mutual respect and empathic understanding
- communication: consistently and honestly presenting not only what is known but what is not known (aka uncertain)

- connection: acknowledging a deeper understanding about and appreciation of the complexities of our world through individual and collective sensemaking, connecting to our intuition, to our community and to the world around us
- commitment: demonstrating professionalism that combines competence and ethics, and which is committed to clear and consistent two-way communication in line with the mechanisms for feedback.

Ensuring there are platforms for civil discourse where all members of a community listen to and learn from the concerns and fears of others contributes to a climate wherein meaningful and productive listening allows diverse and opposing viewpoints to be genuinely heard and discussed.

It is hardly surprising to see a rise in hate speech and efforts to shut down opposing viewpoints in the midst of eroding public trust. One does not need to agree with another's point of view to listen to their concerns. But that listening must be understood as a genuine gesture and not simply a token, box-ticking act or superficial community engagement program. Practising open, honest and consistent communication about reasons for taking specific actions contributes to the transparency about the decisions undertaken for and on behalf of citizens. Furthermore, communication must run both ways – which returns us to the value of listening and seeking out the views of others.

As our cities and technologies grow 'smarter', the collection of personal data becomes more automated and ubiquitous. While many data custodians may set appropriate governance processes in place for holding and using personal data, there is not always a clear plan for engaging the data publics to ensure they feel sufficiently consulted and represented in the data collected from and about them.

Within the wider community, there is also a growing consciousness about the vulnerability of data to misinterpretation, misuse and misappropriation. Indigenous knowledge perspectives and deeply abiding practices of First Nations peoples for grounding in country (locality)²⁷ not only show us a way to pay respect to the land and her inhabitants but also to sensitise us to forms of evidence that extend beyond what might be directly visible at any point in time.

Gaining public trust to work with data

There is a difference between recognising the political need to gain the trust of the people in order to achieve your desired goal and actually believing that a more inclusive and engaged approach to the collection and use of data will improve the validity of the data itself. Building public trust is an ongoing, dynamic process. Social research, as outlined above, is showing us that the more trust there is in any given context, the more risk is accepted. And the more risk-enabled a community, the more resilient it can become. Trust building forges between leaders and a

27 See examples of country-centred design: <https://oldwaysnew.com/>; <https://www.aidr.org.au/media/7760/designing-with-country-discussion-paper.pdf>

community the relationship that is needed to support the innovations and actions conducive to business resilience and performance. We must work from inside out (through training our public servants and analytics professionals) and the outside in (through ongoing community engagement) to keep building and refreshing a government that serves its citizens.

Strategies for gaining trust through action: 'Show how' as well as 'know-how'

Trustworthiness is demonstrated not only through competence ('know-how'), but also through showing how experience, knowledge and training are applied. Transparency, fairness, ethics, accountability and accessibility are all values by which we see organisations seeking to gain public trust. Showing how such values are being applied is far much more powerful than simply stating their importance. Showing how, for instance, data-sharing and privacy-preserving practices are applied to particular situations can contribute to better understanding about and confidence with ways that data is being used.

One means of demonstrating trustworthiness in relation to data is by establishing and communicating data stewardship practices through ongoing reviews. The Australian Data Strategy, for instance, references six principles guiding government management and use of data: accountability, benefits, respect, transparency, protection and use (Department of the Prime Minister and Cabinet 2021). New Zealand's Data Futures Partnership developed a set of non-compulsory guidelines framed around eight key questions about the use of personal data organisations collect. Its project used a process for continual reflection on responses to guide transparency and accountability of data use in relation to:

Value

- What is the data used for?
- What are the benefits and for whom?
- Who will use it?

Protection

- Is the data secure?
- Will data be anonymous?
- Can a person see and correct data about themselves?

Choice

- Will the person be asked for consent?
- Can the person's data be sold or shared? (Data Futures Partnership 2017:6)

The UK Government's Data Ethics Framework (2020) uses three overarching principles (transparency, accountability and fairness) to shape ethical considerations through all stages of a public sector data project. Defining and understanding of

public benefit is the first specific action. The stepped list of further actions in this framework prompt a project team to ensure they are communicating how data is being used, evidencing that use is proportional to the user need, and demonstrating efforts to make data work transparent and accountable.

Globally, trustworthiness is also an important consideration for the deployment of AI. The High-Level Expert Group on Artificial Intelligence (AI HLEG) set up by the European Commission, for instance, published ethics guidelines for trustworthy AI in April 2019 (AI HLEG 2019), intended to guide practitioners toward more ethical and more robust applications of AI. According to these AI HLEG trustworthy AI guidelines, to be trustworthy an AI needs to be:

- lawful – respecting all applicable laws and regulations
- robust – both from a technical and social perspective
- ethical – respecting ethical principles and values. (AI HLEG 2019:5)

As mentioned earlier, it is important to note that such guidelines in and of themselves do not demonstrate the trustworthiness of the AI, which is why further work is then needed to assess their deployment in practice (see, for instance, Zicari et al. 2021, who devise a method for assessing general AI HLEG trustworthy AI guidelines in practice).

Why being 'right' is not enough

Making visible the frameworks that an organisation uses for data protection and data governance needs to acknowledge a core value of democratic, civil society. Citizens can and should expect to have a voice in the decisions being made by their government (and government agencies acting on their behalf) and the way that data practices shape government activities, especially decisions that impact on the everyday life of its citizens.

Co-design frameworks (participatory models and mechanisms for ongoing feedback with sufficiently diverse participants, especially including vulnerable groups) can help address these concerns. However, such frameworks require major paradigm shifts in authority towards more process-based rather than rule-based frameworks and a mindset of inclusive and constant consultation and 'evolving design'.

We must appreciate that 'giving voice' involves explaining rather than telling. Ongoing data advocacy work within the general community is as important as developing specialist practitioners to put ethical data science practices into operation. Building trust in the way we work with data involves continually demonstrating the trustworthiness of our data practices.

Putting principles into practice for now and the near future

We should as a matter of principle be designing systems taking into account the wellbeing of people from whom the data is taken in the first place and which involve them in the process of determining how that is to be done.

Participatory approaches that get the community involved in the design process from start to finish are powerful tools for building trust into any network. Increasingly citizens will expect to be involved in the design of the processes by which data about them is collected and used (see, for example, smart city co-building and ethical test beds involving community input in Cooray et al. 2017 and Riedl 2020). As Jer Thorp (2016) eloquently illustrates, if we want to build data systems as two-way streets that respect the citizens from whom the data is sourced, we should be designing systems that focus on the wellbeing of the very people from whom the data is taken in the first place. If we are to create what Thorp (2016) refers to as 'real, functioning data publics', we need to bring data into public, shared spaces. Public value and public inclusion need to be foregrounded to mitigate the risk of reiterating – or worse still, amplifying – inequities and distrust in the design of government services.

Possible ways forward for community engagement and co-design to demonstrate trustworthiness include:

- pursuing 'interactive accounting' of values of data sharing and privacy to determine:
 - where to draw the line (for example, 'category judgements' for privacy limits, acceptable data use)
 - how to establish 'accountability relations'
- identifying vulnerable populations under-represented or over-represented in any datasets in use through:
 - vulnerability and uncertainty scale discussed earlier
 - embedding reflective practices to remain alert to the missing, misrepresented and under-represented in any analytic work undertaken or planned.

Such deliberate engagements will ensure sectors of the community know when, where and how they can raise concerns over privacy and trust.

Displacement and inequalities can be further magnified in the creation of large-scale networked systems. Participatory and inclusive models of engagement with citizens work because they not only confer trust and legitimacy to the body seeking to gather the data, but those methods of engagement also lead to better, more precise and more usable data.

The fact that the decisions we make, as individuals and as a society, are based on value judgements that are subjective and emotional rather than rational need not be problematic *if* we build and maintain our trust-making frameworks. Ethical data practice may mean learning to make the invisible visible by remaining alert to who (and what) is missing, under-represented or misrepresented in our data practices. It will also mean recognising the value of forming a view of customers as co-learners and co-designers rather than just data points in our planning

Building on the earlier discussion of risk, one of the best ways to navigate uncertainty and risk is through open and honest sharing within trusted relationships that can support your learning and growing.

Governments and leaders have a moral and social obligation to reassure the public about their management of data and analytics processes by using controls, processes and standards, providing greater transparency about the way data is used, and articulating the value of any of the resulting systems and technologies they put in place (see, for example, discussion in OECD 2017 and 2019).

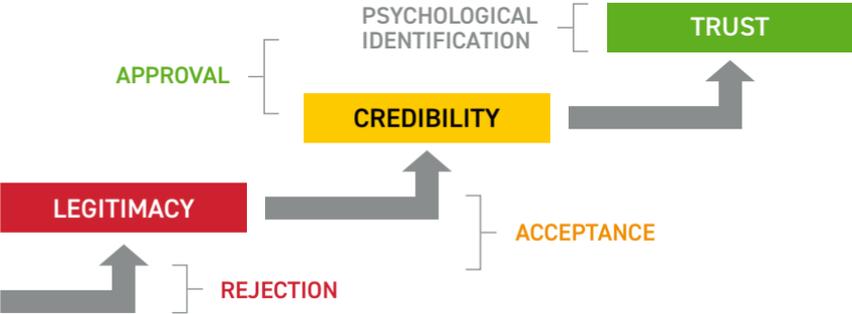
Approaching the governance framework for the data collected by and about the city and her inhabitants in a manner described in the previous section can go some way to building such trust. Overseeing the data on behalf of a community, however, is insufficient on its own.

Demonstrating trust

Trust often takes the form of private or social contracts. Fundamentally, the idea of 'trust' between citizen and state is much more complex than, for example, that between customer and retailer or between friends or colleagues. However, core principles are still the same because building trust involves building a trusted relationship and demonstrating its endurance over time by working within the community and for it. It means accepting risk, responsibility, and accountability for actions.

What it takes to become worthy of trust

Building community acceptance for a particular way of working with data is sometimes referred to as social licence. Thomson and Boutilier (2020) offer a process model derived from social licence initiatives in the resource sector; this might help translate data stewardship practices in terms of a gaining trust and commitment from local communities.



Gaining the social license

Source: Thomson and Boutilier (2020).

Referring to the above diagram, demonstrating trust would involve a stepped approach to building and earning trust. To begin with, legitimacy might be established by taking a human-centred approach to data practices.

Working towards acceptance could involve presenting and explaining, models of appropriate practice throughout the data ecosystem. Credibility could be established by embedding transparency and feedback mechanisms into the deployment and review process.

Trust would be gained and maintained by taking these previous steps and then providing demonstrable outcomes to the community. This process model reminds us that building trust has to happen one step at a time by building from acceptance to approval for actions taken.

Co-designing such a framework involves working from the inside out and the outside in, to keep building and refreshing our practices. This collective effort requires:

- shaping multiperspective foundations intertwining ethical expertise, technical know-how and data practice
- demonstrating to all stakeholders (including the citizens from whom the data is taken in the first place) how competence and ethics shape data sharing practices and policies
- making space for public reflection of the lessons learnt from past practices (successful and not).

Ultimately, true participatory engagement embedded through deployment of co-design processes could thus be considered a way to maintain trust that the organisation has demonstrated and earned.

Building accountable mechanisms of trustworthiness

Good governance is a process of continuous evaluation and communication. There is also great value in framing our approach around a culture of care and stewardship, putting people and ecological flourishing at the centre of everything we do. Framed in this context, devising and explaining criteria for (data) quality can help build trustworthiness into data sharing practices. Much in the same way that social science approaches to data analysis and work in the field help demonstrate the trustworthiness of data collected, we can build trustworthiness into our data practices through:

- prolonged engagement in the problem space
- persistent observation to gain rich insight and understanding of context
- triangulation to ensure multiple points of data shape analysis
- referential adequacy checks to bake quality criteria into analytic practice
- peer debriefing to bring diverse expertise and feedback into interpretation of data
- regular check-ins with community members to stay closely connected to data sources and community perspectives on data interpretation and project value.

Documenting these field processes in use makes visible the ways in which trustworthiness can be judged and evaluated. It lays bare the way that decisions are

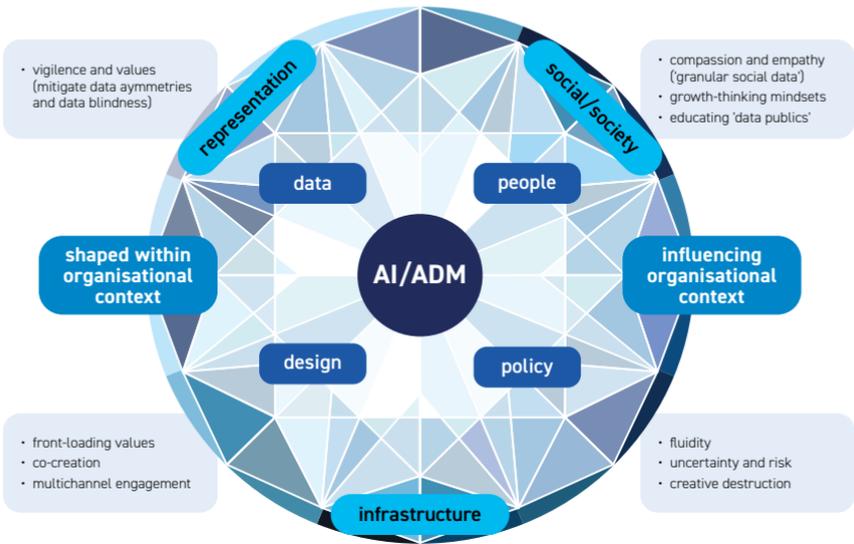
made – even in light of uncertain situations and incomplete information – allowing pathways for feedback and engagement with the process.

Audits built on these principles contribute to trust-building and trust-preserving activities. This is more than just a way to build public trust; it's also a way to build trust in the data claims themselves. Consequently, criteria of data trustworthiness and data soundness can contribute to overall data quality.

A foundational framing for building trusted partnerships for AI and data

This section of the chapter demonstrates a framework that applies these key understandings about trust building specifically to data and AI technologies.

Appreciating the landscape of knowledge production in all its entangled human-machine complexity is vital first step to taking a participatory approach to building ethical, evidence-based decision-making frameworks making use of citizen data. The working model shown below provides a sociotechnical framework for ethical data practice drawing attention to four key interwoven components: data, design, people and policy.



Locating trust building and ethical data practice in a sociotechnical framework

Source: Anderson, v2, November 2019.

*ADM = automated decision-making

The fractal representation of these four components of the framework is intended to act as an analytic aid to engage with the background 'shadow work' (Sawyer and Tapia 2006; Star and Strauss 1999) of data practices within any entity's organisational and political contexts. It draws attention to ways integration of participatory models and mechanisms for ongoing feedback with the community can contribute to building, demonstrating and sustaining trustworthiness.

Processes on the left-hand side of the diagram are more within the control of an organisation. Here is where trust-building mechanisms for working with *data* can be deliberately designed. There is increasing recognition that governments and organisations need to provide greater transparency about the way data is collected and used (for example, AI HLEG 2019; OECD 2017 and 2019). Efforts to reassure the public about the management of data and analytics using controls, processes and standards, for instance, are part of this *design* process.

The right-hand side of the diagram portrays features of this sociotechnical context that are far less in our direct control: the *people* and *policy* components of the framework. The assemblages that evolve in this human-machine-information interplay rarely lend themselves to deliberate design and yet, ironically, as they become more naturalised and more invisible, their configuration can become more frozen.

A growing number of cases demonstrate how algorithms have replicated or even exacerbated inequalities in the ways that different demographic groups are treated. What's more, if left unchecked, the assumptions and values embedded in these technologies and the decisions they enable can become baked into the infrastructures that drive subsequent knowledge practices.

There is also a growing cry to 'turn data around' and design data systems that take into account the wellbeing of people from whom the data is taken in the first place (Thorp 2016). Data does not speak for itself but, rather, is given a voice by the people and the algorithms that play increasingly critical roles in the transformation of data into insight.

Public trust is not lightly given – it is earned over time and is an ongoing process of engagement with the community. Recently studies finding that the raging infodemic is feeding mistrust are also finding that globally citizens are beginning to appreciate the importance of information and science literacies, political awareness and speaking out when there is a need for change and reform (Edelman 2021:23–25; Ienca and Vayena 2020; OECD 2022). In an interview given in the early months of the COVID-19 pandemic, Bruno Latour observed:

If you want people to have some grasp of science, you must show how it is produced. (Watts 2020)

Genuine engagement, partnered with explaining the how and why of our actions can contribute reassurance. Thus, deliberately and frequently engaging with the 'data publics' represented in any datasets in use can be critical for evaluating the effectiveness of any trust-building work undertaken.

The fractal framing of data's sociotechnical intertwinings presented here is intended to help push past binarism and appreciate the impact of data representations and challenges of categorisation in relation to our data practice.

Raising awareness of the social and political factors outside the control of an organisation can draw attention to ways assumptions about the value of attributes in a dataset can perpetuate prejudice and inequity, with those marginalised becoming further disenfranchised.

There is a value-laden chain of activities inviting more thoughtful consideration about who/how/what is counted and analysed in these increasingly data-intensive spaces. Tools are made through practice. In helping policymakers and practitioners to appreciate ways that human judgements and values can and should be permitted to augment the computational components of the data assemblages they are shaping, the human capacity to work with uncertainty and intuitive judgements comes to be seen as an essential partner to high-powered computational and analytical capacities.

Ways forward – Building public trust with wellbeing as our driver

Good governance frameworks steward data assets and oversee outcomes in line with the core values of the community. With wellbeing as the driver, such a framework provides the assurances of safety and security necessary to enable a community to sit (more) comfortable in uncertainties and ultimately to flourish.

It is therefore critically important that the welfare of the most vulnerable members of any community are looked after and that multistakeholder perspectives figure in the governance of any data deployments. The principles of Indigenous Data Governance (IDG) provide an excellent model for all data governance by alerting us to the powerful controls that data can exert on the most vulnerable sectors of a community. As Carroll et al. (2019) articulate in their exploration of IDG:

Indigenous data governance can thus be described as a reciprocal relationship between data for governance and governance of data. The first is a matter of quality, relevance, and access: can Native nations obtain the data they need for governance? The second is a matter of ownership and control: can Native nations manage, protect, and use that data?

Operating principles drawn from IDG can also provide us with guidance about ways to demonstrate trustworthiness to the public in relation to the way that social data is used, most notably by offering guidance for navigating the complexities of individual and collective rights:

Through this communal lens, Indigenous peoples conceptualize IDS not only as a right, but also as a responsibility. (Carroll et al. 2019)

Following on from this, trustworthiness would be demonstrated by focusing on relationships, the wellbeing of the community, and responsibilities for stewarding for future generations.

If wellbeing is allowed to determine the ethics of a system, Carroll et al. (2019) argue, governance is in the service of a community's 'foundational capacity' to make and implement strategic decisions about their own affairs. Embedding trust-building activities in all aspects of data practice guides the creation of such a system. Putting this trust-building work into the four quadrants discussed above might look as follows:

- reassurance: repairing trust deficit using markers of democratic process (communication, consultation)
- resilience: nurturing capacity for insight and innovation in face of challenges, supporting diversity, encouraging curiosity, enabling flourishing
- relationships: partnering to build trust essential to tackle complex, interconnected challenges and putting expertise into context
- reflection: making time to think and learn from success and failure; supporting education and empathic understanding.

Engaging with trust in the way discussed in this section enriches our understanding of what it means to act in the public interest. Thinking in terms of these quadrants, building of trust is not a gate to pass through on the way to completing a project, but rather becomes a core function of our work that can not only ensure transparency in decision-making, but also continuously connect us to community and context.

Understanding the trust deficit – What we need to learn to move forward

As has been emphasised in this chapter, the judgement of trust is a qualified judgement, not a rational one. While it is not the only mechanism influencing the relations between citizens and the state, analysis of past events (especially epidemics and moments of crisis) underscores how the building of relationships ahead of time can give leaders and experts a crucial edge when seeking to gain the public's trust in times of crisis. These events also shed light on the impact that local knowledge of and engagement with the community can have on the perception of trustworthiness.

Why decision-makers must take the work of trust building seriously

Trust buys time to work through uncertainties, doubts and crises. As Luhmann observes in his reflection on the relation between trust and power:

Through trust a system gains time, and time is the critical variable in the construction of complex system structures. The satisfying of needs can be delayed, and nevertheless guaranteed. Instrumental action, oriented towards distant effects, can become institutionalised if the temporal horizon of a system is suitably extended by means of trust. The availability of liquid financial resources, power and truth, all mechanisms dependent on trust, makes possible an indifference on the part of the system towards numerous events in the environment and thus a gain in reaction time. (Luhmann 2018:98)

The confidence that is demonstrated (and earned) through the building of a trusted relationship is critical if we, as individuals, organisations and as leaders, are to help our communities to accept the challenges of working with the complexities of any given situation.

In their research into the responsible use of data to tackle the pandemic, Ienca and Vayena (2020) discuss the implications of the link between mistrust in COVID-19 reporting and low levels of trust in government:

This risk of mistrust is even greater in countries in which citizens place a much lower level of trust in their government, such as Italy, France and the USA. Therefore, whenever access to these data sources is required and is deemed proportional, the public should be adequately informed. Secrecy about data access and use should be avoided. Transparent public communication about data processing for the common good should be pursued. Data-processing agreements, for example, should disclose which data are transmitted to third parties and for which purpose. (Ienca and Vayena 2020:464)

Local and international studies report that societal leaders are not trusted to handle challenges. Results from a Pew Research Center study of the attitudes of US adults reported in September 2020 found the share of Americans (of either political persuasion) who say they trust the federal government to do what is right either just about always or most of the time has hovered near 20% since the global financial crisis of 2008.

This contrasts with the Lyndon B Johnson administration, when more than 73% of Republicans and 80% of Democrats trusted Washington always or most of the time (Pew Research 2020:14). An Australian election study yielded similar results, with 25% of respondents in the 2019 survey stating they believe people in government can be trusted, compared to 51% reported in the 1969 survey (Cameron and McAllister 2019:99).

These national findings in the US and Australia also correspond with the international results reported in Edelman Trust Barometer reports of the past three years, where having confidence in societal leaders 'to do what is right' continues to decline. Furthermore, according to the 2021 findings of the Edelman Trust Barometer, the infodemic accompanying the COVID-19 pandemic pushes people to trust local sources more than government or institutional ones (Edelman 2021:23–25).

Growing disparities in the digital economy continue to breed disenfranchisement and disconnection, which in turn contributes to a trust deficit rife with disinformation, conspiracy theories and disaffection. As social capital fragments, trust breaks, and people return to what and who they know they can trust: those closest to them.

A key takeaway from this trust deficit is this: regardless of how cutting-edge and valid expertise or datasets may be, without a foundation of trust and good relationships with our peers and the communities we hope to serve, it can prove difficult for others to accept new knowledge and change their behaviour.

Moving forward

Even in an age of big data and evidence-based decision-making, trust building remains as personal, local and ultimately political as it ever was. To be worthy of trust takes more than authority. The good governance associated with demonstrating trustworthiness in unusual times requires better listening.

Trust in a democratic society relies on genuine two-way communication with communities and individuals. Being accountable to the community involves not only communicating decisions and actions clearly and consistently, but also ensuring that concerns of the community (particularly those most vulnerable) are genuinely heard.

If we are to learn lessons from the trust deficit witnessed in recent years, we need to ensure that:

1. We recognise that both competence and ethics must be intertwined and demonstrated to the community.
2. We appreciate that public trust is not lightly given, but rather earned over time in an ongoing process of engagement with the community. Building trusted relationships involves demonstrating at the local level the ways that government is accountable (even when something does not go according to plan), capable and credible. It shows that government is designing with the community and not simply for them. Most importantly, its ongoing nature means that it is already in place in times of crisis, and this is particularly important in relation to vulnerable populations.
3. We value co-designing privacy policies as an essential component for building public trust. Giving voice is embedded in the fabric of the democratic values we strive to put into practice as a community. Not only is it essential, but it also has powerful benefits. Taking seriously the democratic, participatory ideals held so dear by our society has implications for our practice in relation to data and privacy.
4. We make space for public reflection of the lessons learnt – demonstrating what has been learnt from past practices, successful or not. It involves making time to pause and reflect on lessons from past and present for the future. Investing the time and effort to build a trusted relationship buys the time necessary to navigate the inevitable uncertainties of a world where information (and data) will always be imperfect and incomplete.
5. We invest in the critical work of public and professional education and training. In a democratic society, trust involves trusting the capacity of the people (people are expert in their own contexts/communities/individual situations). We need to take seriously the responsibility to create better educated 'data publics' as part of the co-design process.

In times of crisis and high uncertainty, trust becomes more tenuous even as it becomes more necessary. The need to demonstrate trustworthiness increases in line with the

sense of vulnerability of the community (or sector of that population). For this reason, if regimes of trust-building have been established prior to crisis, the community is more likely to trust the actions of government or others acting in their presumed public interest. Building trust prior to crisis supports the fluidity needed during crisis.

References

- Anderson TD (2006) 'Uncertainty in Action: Observing Information Seeking Within the Creative Processes of Scholarly Research', *Information Research*, 12(1), paper 283. <http://InformationR.net/ir/12-1/paper283.html>
- Anderson TD (2013) 'The 4Ps of Innovation Culture: Conceptions of Creatively Engaging with Information', *Information Research*, 18(3), paper C28. <http://InformationR.net/ir/18-3/colis/paperC28.html>
- Anderson TD (2020) 'Keystone Practices to Enable Smart Cities to Flourish' in *Geography Research Forum*, 40(1):171–192. <https://grf.bgu.ac.il/index.php/GRF/article/view/603>
- Barfield O (1993) *A Barfield Sampler: Poetry and Fiction* by Owen Barfield (Clayton Hunter J and Kranidas T eds), SUNY Press.
- Beck U (1992) *Risk Society: Towards a New Modernity*, Sage Publications.
- Beck U (2000) 'Risk Society Revisited: Theory, Politics and Research Programmes' in Adam B, Beck U and van Loon J (eds) *The Risk Society and Beyond: Critical Issues for Social Theory*, Sage Publications.
- Boholm Å (2003) 'The Cultural Nature of Risk: Can There Be an Anthropology of Uncertainty?', *Ethnos*, 68(2):159–178.
- Burke K (24 November 2020) 'Overcoming Mythinformation', *Australian Pharmacist*. <https://www.australianpharmacist.com.au/overcoming-mythinformation/>
- Cameron S and McAllister I (December 2019) *Australian Election Study 1987–2019*. <https://australianelectionstudy.org>
- Carroll SR, Rodriguez-Lonebear D and Martinez A (2019) 'Indigenous Data Governance: Strategies from United States Native Nations', *Data Science Journal*, 18(1):31. <https://doi.org/10.5334/dsj-2019-031>
- Christensen P and Mikkelsen MR (2008) 'Jumping Off and Being Careful: Children's Strategies of Risk Management in Everyday Life', *Sociology of Health and Illness*, 30(1):112–130.
- Cooray M, Duus R and Bundgaard L (23 August 2017) 'Technology is not Enough to Create Connected Cities – Here's Why', *The Conversation*. <https://theconversation.com/technology-is-not-enough-to-create-connected-cities-heres-why-82740>
- Craven M, Mysore M, Singhal S and Wilson M (13 April 2020) 'COVID-19: Briefing note #5', *COVID-19: Implications for business in 2020*, McKinsey & Company. <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/covid-19-implications-for-business-2020>
- Data Futures Partnership (2017) *A Path to Social Licence: Guidelines for Trusted Data Use*. https://aisp.upenn.edu/wp-content/uploads/2019/08/Trusted-Data-Use_2017.pdf
- Department of the Prime Minister and Cabinet (2021) *Australian Data Strategy: The Australian Government's Whole-of-Economy Vision for Data*. <https://ausdatastrategy.pmc.gov.au/sites/default/files/2021-12/australian-data-strategy.pdf>

- Edelman (n.d.) Edelman Trust Barometer Archive. <https://www.edelman.com/trust/archive>
- Edelman (2021) Edelman Trust Barometer 2021. <https://www.edelman.com/sites/g/files/aatuss191/files/2021-03/2021%20Edelman%20Trust%20Barometer.pdf>
- Edelman R (n.d.) 20 Years of Trust. <https://www.edelman.com/20yearsoftrust/>
- Edelman R (2020) The Evolution of Trust. <https://www.edelman.com/research/evolution-trust>
- Eisenberg EM (2001) 'Building a mystery: Toward a new theory of communication and identity', *Journal of Communication*, 51(3):534–552.
- Giddens A (1990) *The consequences of Modernity*, Stanford University Press.
- High-Level Expert Group on Artificial Intelligence (2019) *Ethics Guidelines for Trustworthy AI*, European Commission. <https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html>
- Huff M and Rea P (2009) Deconstructing Deceit: 9/11, the Media, and Myth Information. <http://www.projectcensored.org/wp-content/uploads/2010/05/DeconstructingDeceitOnlineEd.pdf>
- Ienca M and Vayena E (2020) 'On the responsible use of digital data to tackle the COVID-19 pandemic', *Nature Medicine*, 26:463–464. <https://doi.org/10.1038/s41591-020-0832-5>
- Jaffe D (2018) 'The Essential Importance of Trust: How to Build It or Restore It', *Forbes*. <https://www.forbes.com/sites/dennisjaffe/2018/12/05/the-essential-importance-of-trust-how-to-build-it-or-restore-it/?sh=68f3eb0c64fe>
- KPMG (2018) *Guardians of Trust: Who is Responsible for Trusted Analytics in the Digital Age?* <https://assets.kpmg/content/dam/kpmg/xx/pdf/2018/02/guardians-of-trust.pdf>
- Luhmann N (2018) *Trust and power*. John Wiley and Sons.
- Maguen S, Papa A and Litz BT (2008) 'Coping with the threat of terrorism: A review', *Anxiety, Stress, & Coping*, 21(1):15–35. <https://doi.org/10.1080/10615800701652777>
- Malaby TM (2002) 'Odds and Ends: Risk, Mortality, and the Politics of Contingency', *Culture, Medicine, and Psychiatry*, 26(3):283–312.
- Minors D (2021) 'Understanding the Infodemic', *Curiosity (Research Magazine of the University of the Witwatersrand, Johannesburg)*, 11:6–27.
- Neyland D (2006) *Privacy, Surveillance and Public Trust*. Springer.
- Organisation for Economic Cooperation and Development (2017) *Government at a Glance 2017*, OECD Publishing. https://doi.org/10.1787/gov_glance-2017-en
- Organisation for Economic Cooperation and Development (2019) *Embracing Innovation in Government: Global Trends 2019*. <https://www.oecd.org/innovation/innovative-government/embracing-innovation-in-government-global-trends-2019.htm>
- Organisation for Economic Cooperation and Development (2022) 'Building Trust to Reinforce Democracy: Main Findings from the 2021 OECD Survey on Drivers of Trust in Public Institutions', *Building Trust in Public Institutions*, OECD Publishing. <https://doi.org/10.1787/b407f99c-en>
- Pew Research Center (14 September 2020) 'Americans' Views of Government: Low Trust, but Some Positive Performance Ratings'. <https://www.pewresearch.org/politics/2020/09/14/americans-views-of-government-low-trust-but-some-positive-performance-ratings/>

- Riedl DL (June 2020) 'Toward Inclusive Urban Technology', Benton Institute for Broadband & Society. <https://www.benton.org/publications/inclusive-urban-tech>
- Rosa EA (1998) 'Metatheoretical Foundations for Post-Normal Risk', *Journal of Risk Research*, 1(1):15–44.
- Sawyer S and Tapia A (2006) 'Always Articulating: Theorizing on Mobile and Wireless Technologies', *The Information Society*, 22(5):311–323. <http://doi.org/10.1080/01972240600904258>
- Star SL and Strauss A (1999) 'Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work', *Computer Supported Cooperative Work (CSCW)*, 8(1):9–30. <https://doi.org/10.1023/A%3A1008651105359>
- Taddeo M (2009) 'Defining Trust and E-Trust: From Old Theories to New Problems', *International Journal of Technology and Human Interaction*, 5(2):23–35.
- Thomson I and Bouillier R (2020) What is the Social License? <https://sociallicense.com/definition.html>
- Thorp J (2016) 'Turning data around', *Office for Creative Research Journal*, 2:11–24. (Also available at <https://medium.com/memo-random/turning-data-around-7acea1f7479c#tf2wc1p8e>)
- Tulchinsky TH (2018) Case Studies in Public Health. <http://doi.org/10.1016/B978-0-12-804571-8.00017-2>
- UK Government (2020) Data Ethics Framework. <https://www.gov.uk/government/publications/data-ethics-framework#the-data-ethics-framework-principles>
- Uslaner EM (2003) 'Varieties of trust', *European Political Science*, 2(3):43–49.
- Uslaner EM and Brown M (2005) 'Inequality, Trust, and Civic Engagement', *American Politics Research*, 33(6):868–894.
- Wallerstein I (1998) 'Uncertainty and creativity', *The American Behavioral Scientist*, 42(3):320–322.
- Watts J (2020) 'Bruno Latour: "This is a global catastrophe that has come from within"', *The Guardian*. <https://www.theguardian.com/world/2020/jun/06/bruno-latour-coronavirus-gaia-hypothesis-climate-crisis>
- World Health Organisation (2009) WHO Guidelines on Hand Hygiene in Health Care. <https://www.ncbi.nlm.nih.gov/books/NBK144018/>
- Wilson TD, Centerbar DB, Kermer DA and Gilbert DT (2005) 'The Pleasures of Uncertainty: Prolonging Positive Moods in Ways People Do Not Anticipate', *Journal of Personality and Social Psychology*, 88(1):5–21. <http://doi.org/10.1037/0022-3514.88.1.5>
- Winner L (1984) 'Mythinformation in the high-tech era', *IEEE Spectrum*, 21(6):90–96.
- Zaloom C (2004) 'The productive life of risk', *Cultural Anthropology*, 19(3):365–391.
- Zicari RV, Brodersen J, Brusseau J, Dudder B, Eichhorn T, Ivanov T, Kararigas G, Kringen P, McCullough M, Möslin F, Mushtaq N, Roig G, Stürtz N, Tolle K, Tihi JJ, van Halem I and Westerlund M (2021) 'Z-Inspection®: A Process to Assess Trustworthy AI', *IEEE Transactions on Technology and Society*, 2(2): 83–97. <https://doi.org/10.1109/TTS.2021.3066209>

Chapter 5

Data privacy, fairness and privacy harms in an algorithm- and AI-enabled world



By Peter Leonard

Peter Leonard is a data and technology consultant and lawyer advising data-driven businesses and government agencies. Peter is principal of Data Synergies and a Professor of Practice at UNSW Business School (IT Systems and Management, and Management and Governance). Peter also serves on the NSW Government's AI Review Committee, tasked to "guide and provide strategic oversight for the use of AI in [NSW] government", and the NSW Information and Privacy Advisory Committee, tasked to "provide the government with information, advice, assistance and training to deliver world-leading information and privacy management practices".

Slumberous autumn when a little dog barks, runs before rain descends

There is growing consensus that the Australian Privacy Act, in common with similar statutes in other jurisdictions, needs a major overhaul.

The review by the Australian Attorney-General's Department (AGD) of the *Privacy Act 1988* (Cth) (Privacy Act),²⁸ underway at the time of writing this piece, provides an opportunity to:

- improve mechanisms to protect data privacy of Australians
- reduce friction of cross-border dealings, by improving alignment of Australian data privacy regulation with international regulatory best practice
- accommodate societally beneficial secondary and derived uses of data.

Australia has the opportunity to select and tailor the best features of new data privacy statutes from around the world and to ensure that the Privacy Act belatedly becomes fit for purpose in the 21st century.

There are many current initiatives for reform of data privacy laws in comparable jurisdictions that should inform overhaul of the Australian Privacy Act. They include:

- a comprehensive review in the UK of whether UK GDPR should diverge from EU GDPR, with the stated objective of better enabling innovation in the UK
- proposals in the European Union to supplement EU GDPR with a Digital Markets Act and a Digital Services Act, and an associated package of initiatives to address applications of AI and advanced data analytics²⁹
- in the US, proposals for a federal data privacy statute,³⁰ development by the Uniform Law Commission of a Uniform Personal Data Protection Act,³¹ and US state by state enactment of data privacy statutes³²
- substantial recent revisions of data privacy statutes in Singapore, Korea and Japan, and a new statute in Quebec

28 The Australian Privacy Act is available at <https://www.legislation.gov.au/Details/C2021C00452>

29 For an analysis of the interaction between proposed provisions of the Digital Markets Act and GDPR, see the Centre for Information Policy Leadership Bridging the DMA and the GDPR, December 2021, at <https://www.huntonprivacyblog.com/2021/12/16/cipl-publishes-white-paper-on-the-interplay-between-the-draft-eu-digital-markets-act-and-the-gdpr/>

30 For a summary, see IAPP, US Federal Privacy Legislation Tracker, <https://iapp.org/resources/article/us-federal-privacy-legislation-tracker/>

31 Uniform Personal Data Protection Act, as drafted by the US Uniform Law Commission, is linked at <https://fpf.org/blog/uniform-law-commission-finalizes-model-state-privacy-law/>.

32 See US State Privacy Legislation Tracker, <https://iapp.org/resources/article/us-state-privacy-legislation-tracker/>

- proposed revisions to the Canadian federal privacy statute and for a new data protection statute in India.

Many consumer organisations and privacy advocates across the world criticise national privacy and data protection statutes, and enforcement of them, as inadequate and incomplete. Sometimes those criticisms are echoed within international organisations and national legislatures (United Nations High Commissioner for Human Rights 2021). As stated by the UK Parliament's House of Commons and House of Lords Joint Committee on Human Rights (2021) in its Inquiry Report on *The Right to Privacy (Article 8) and the Digital Revolution*:

The evidence we heard during this inquiry ... has convinced us that the consent model is broken. The information providing the details of what we are consenting to is too complicated for the vast majority of people to understand. Far too often, the use of a service or website is conditional on consent being given: the choice is between full consent or not being able to use the website or service. This raises questions over how meaningful this consent can ever really be.

While most of us are probably unaware of who we have consented to share our information with and what we have agreed that they can do with it, this is undoubtedly doubly true for children. The law allows children aged 13 and over to give their own consent. If adults struggle to understand complex consent agreements, how do we expect our children to give informed consent? Parents have no say over or knowledge of the data their children are sharing and with whom. There is no effective mechanism for a company to determine the age of a person providing consent. In reality a child of any age can click a 'consent' button.

The bogus reliance on 'consent' is in clear conflict with our right to privacy. The consent model relies on us, as individuals, to understand, take decisions, and be responsible for how our data is used. But we heard that it is difficult, if not nearly impossible, for people to find out whom their data has been shared with, to stop it being shared or to delete inaccurate information about themselves. Even when consent is given, all too often the limit of that consent is not respected. We believe companies must make it much easier for us to understand how our data is used and shared. They must make it easier for us to 'opt out' of some or all of our data being used. More fundamentally, however, the onus should not be on us to ensure our data is used appropriately – the system should be designed so that we are protected without requiring us to understand and to police whether our freedoms are being protected.

As one witness to our inquiry said, when we enter a building we expect it to be safe. We are not expected to examine and understand all the paperwork and then tick a box that lets the companies involved 'off the hook'. It is the job of the law, the regulatory system and of regulators to ensure that the appropriate standards have been met to keep us from harm and ensure our safe passage. We do not believe the internet should be any different. The Government must

ensure that there is robust regulation over how our data can be collected and used, and that regulation must be stringently enforced.

Notwithstanding such concerns, reform of data privacy law in various jurisdictions is slow and highly contested. This chapter considers why this is the case.

We then explore some key issues enlivening debate on the appropriate scope of reform of Australian data privacy law, with a particular focus upon proposals for reform of the Privacy Act and comparable state and territory data privacy and health information statutes.

We then review the role for data privacy impact assessment in improving accountability of regulated entities for their acts and practices affecting data privacy. Existing practices in data privacy impact assessment are of variable quality. We examine why this is the case and how this should lead to concern that proposed tools for AI and algorithmic impact assessment may not be properly developed and reliably applied by the broad range of entities already deploying and using automated decision-making.

This chapter concludes with an opinionated design manifesto for reform of the Australian Privacy Act, aimed at ensuring the statute becomes fit for purpose for the 21st century – albeit at more than two decades into that century.

Building good statutes requires good policy foundations. We start by asking two foundational questions:

- What should a data privacy statute do?
- What should a data privacy statute not do?

We also caution that Australian data privacy regulation should align with international best practice. Many entities regulated under Australian data privacy laws already conduct operations in multiple jurisdictions or have ambitions to do so. If Australia elects to chart its own course, Australian entities may be forced to incur substantial regulation-induced costs in adapting data architectures and analytics processes, and data handling practices, for cross-border dealings.

In any event, there are emerging convergences in key settings in data privacy statutes in Australia, New Zealand, Singapore, Japan and Korea, most notably in relation to settings around use of privacy-enhancing technologies and controlled data analytics environments that rely upon effective anonymisation. These convergences provide opportunities for further alignment and friction-reducing measures, such as mutual recognition schemes across those jurisdictions.

Australian policymakers should exercise particular caution to avoid, wherever reasonably practicable, devising regulatory measures that lead to Australia-specific regulation-induced costs for Australian entities in cross-border dealings.

Concerns as to the collection and uses of data about consumers and the scope of data privacy law

Data policy concerns now range far beyond the scope of rights or interests of citizens to go about their private lives, including in public and semipublic places, without unjustified or unexpected collection and uses of data. The range of concerns as to the collection and uses of data about consumers and other citizens continues to grow, and includes:

- the relative roles of consideration by regulated entities of social responsibility, business ethics or social licence to moderate and control unjustified or unexpected collection and uses of data, and enactment and enforcement of 'hard law' with penalties and legal sanctions
- the need to nurture digital trust of citizens in order to ensure a vibrant digital economy
- the importance of digital inclusion and addressing accessibility of digital services by all
- the facilitation of societally beneficial uses of data³³
- considerations of social equity, and the reasonableness (or otherwise) of weighting of benefits for the many against detriments to a few
- online safety and protection of children and other vulnerable people
- focus on, for online services, use by service providers of 'dark patterns' and behavioural psychology to encourage individuals to volunteer data or to not seek out privacy options and exercise them to shift settings to be more privacy protective
- attention to the emerging panopticon of surveillance and 'profiling' of citizens
- 'biased and discriminatory' algorithms and AI
- 'unaccountable' algorithms and AI
- lack of transparency of privacy intrusive acts and practices of businesses, governments, political parties and other political actors, and some not-for-profits
- limitations in legal authority and practical ability of national actors and national regulation to address cross-border and global issues, including acts and practices of entities operating in other jurisdictions and dealing from outside the jurisdiction with citizens or residents within the jurisdiction
- considerations of national political and economic sovereignty and protection from foreign political interference
- addressing growing capabilities of hackers and other malicious actors to exfiltrate sensitive data about consumers and other citizens, and to disrupt supply chains and food security

33 See further, Chan J and Saunders P (2021) Big Data for Australian Social Policy, <https://socialsciences.org.au/wp-content/uploads/2021/12/Big-Data-for-Australian-Social-Policy.pdf>

- whether, how and for which industry sectors, to facilitate portability of consumer data as a tool to empower consumers to compare offerings and switch between providers of products or services and thereby facilitate disruption of incumbents
- whether or when to protect and promote 'national champions' against offshore service providers, including global digital platforms, or to otherwise use consumer data as a tool in 'industry policy' regulation to effect structural adjustments within a national economy.

Sometimes it is not even clear which of the above concerns, or whether other concerns, belong to a debate about data policy settings, or who needs to be engaged as relevant stakeholders to properly inform a debate.

Data policy debates are therefore no longer just about privacy, or principally about data derived from online activity of internet users.

The continuing relevance of and need for data privacy statutes and privacy-focused regulators does not appear to be seriously contested. However, because operations of businesses, governments and other organisations are increasingly enabled by applications of advanced data analytics and AI, in many jurisdictions national policymakers are actively considering adjustments in the relative roles and functions of consumer protection, competition (antitrust) and data privacy (protection) regulators.

Some jurisdictions have proposed re-siting of regulatory responsibilities in relation to data privacy. Competition (antitrust) and consumer protection statutes, and the regulators enforcing them, have steadily gained significant status relative to data protection (privacy) statutes. This trend is in part due to primary reliance by regulators upon provisions in competition statutes, or consumer protection statutes, to address policy concerns as to data handling practices of large online platforms and social media networks.³⁴ Competition powers are now often being used to require large online platforms to implement non-structural safeguards, including operational separation and accountability measures, under the supervision of competition regulators, and not data protection regulators.³⁵

34 For example, in the UK, A new pro-competition regime for digital markets, consultation paper for UK Parliament, CP 489, July 2021; in the US, David N. Cicilline (RI-01) and Ken Buck (CO-04), House Lawmakers Release Anti-Monopoly Agenda for "A Stronger Online Economy: Opportunity, Innovation Choice, media release of 11 June 2021 and accompanying bills as linked in that release; House Judiciary Committee, Judiciary Antitrust Subcommittee Investigation Reveals Digital Economy Highly Concentrated, Impacted By Monopoly Power, media release of 6 October 2020 and the report (Investigation of Competition in the Digital Marketplace: Majority Staff Report and Recommendations).

35 See, for example, Rod Sims, ACCC Chair, paper entitled Competition in Australia faces big challenges delivered to the UniSA and ACCC Competition Law and Economics Workshop, 15 October 2021, at <https://www.accc.gov.au/speech/competition-in-australia-faces-big-challenges>.

Protecting some interests of individuals in data privacy

The Australian *Privacy Act 1988* is misleadingly labelled. The Act does not confer a legal right of individuals in and to data privacy. The Act addresses only a subset of the set of rights of privacy of individuals as commonly asserted and as referred to in international conventions and declarations of human rights.³⁶ The Privacy Act could be more accurately described as the *Data Privacy Act*, where legal protection of data privacy interests of citizens is intermediated by the Australian Information Commissioner.

The Australian Privacy Act is intended to empower individuals by informing them how data about them may be being collected, used and disclosed, and thereby enable them to exercise a choice. The mechanisms to give effect to these objects are variously called 'notice and consent', 'notice and choice', 'individual choice' or 'privacy self-management'. The underlying theory is that an affected individual is afforded 'transparency' as to privacy-affecting acts and practices of a regulated entity, and may then make a choice about whether to deal with that entity. The statute:

- provides a framework of legal principles that regulated entities are required to comply with regarding permitted acts and practices in collecting and dealing with information about identified or identifiable individuals
- specifies when and how affected individuals must be informed how data about them may be being collected, used and disclosed.

The Act has limited coverage. Significant sectors of the Australian economy are exempted, including small business, politicians and political parties, media when conducting journalism, persons acting in a personal or domestic capacity, and state and territory agencies.

Restrictions within the Privacy Act are overridden to the extent a particular act or practice is required by or under an Australia law or a court/tribunal order.³⁷ Legal compulsion under any other federal, state or territory statute, or by subpoena or other court order, prevails over restrictions in the Privacy Act. The Privacy Act does not require a regulated entity, or an authority compelling a disclosure, to weigh reasonable proportionality of the legal compulsion against interests of an affected individual in their data privacy.

Some empowering statutes require weighing by an authority of proportionality, or other consideration of balancing factors. Many empowering statutes do not. Many empowering statutes also do not require independent review, do not require review by senior management, or provide judicial consideration of whether to exercise a proposed legal compulsion. A regulated entity is not required to consult with an affected individual before a disclosure, even where the relevant disclosure would not

36 See the discussion of human rights law in Australia in Australian Human Rights Commission, *Human Rights and Technology Final Report*, March 2021 at <https://humanrights.gov.au/our-work/rights-and-freedoms/publications/human-rights-and-technology-final-report-2021>

37 A number of the Australian Privacy Principles (APPs) provide an exception if an APP entity is 'required or authorised by or under an Australian law or a court/tribunal order' to act differently.

prejudice investigations or other activities of law enforcement agencies or national security organisations.

The Australian Privacy Act, and similar state and territory statutes (which address privacy-affecting activities of state and territory government agencies, local government and some private sector providers of health services), address collection and handling of data about identifiable individuals but not privacy harms that may arise from intrusive and excessive deployment and use of surveillance technologies and geo-tracking devices. A variety of inconsistent state and territory statutes provide some protections in relation to use of surveillance and tracking devices. Surveillance technologies and geo-tracking devices may capture data about identifiable individuals that is then a collection of personal information regulated by the relevant data privacy statute.

A right to know, complain, and elect not to deal: Not a legal right of privacy

Accordingly, the privacy statutes address collection and uses of data about individuals, not broader protection of privacy.³⁸ Instead of conferring a legally enforceable right of individuals in and to data privacy, the Act states as its first two 'objects':

- 'to promote the protection of privacy of individuals'
- 'to recognise that the protection of privacy of individuals is balanced with the interests of entities in carrying out their functions or activities'.³⁹

The Act principally gives effect to these objects by requiring affected individuals to be informed how personally identifying data about them will be collected, used and disclosed, and the purpose for which this will occur. Each regulated entity is required to assess the 'reasonable necessity' of that act of practice to achieve that stated purpose, and to 'balance' its self-interest in collecting and using that data with that entity's assessment of expectations of different sections of the public in 'protection of privacy' and the extent to which those expectations are fair and reasonable.

Consistent with this regulatory theory, if data as collected is non-identifying, or to be used has been transformed so that the data and outputs from analysis of that data is reliably and pervasively de-identified (effectively anonymised), the Act does not operate in relation to uses and disclosure of that effectively anonymised data. That noted, effectively anonymised data may still enable differentiation in treatment between unidentifiable individuals based upon inferences as to activities, interests, preferences and characteristics of those and other ('like') unidentifiable individuals.

38 Most data privacy statutes do not define 'privacy' and there is a surprising diversity of definitions of 'privacy'. See further, Cohen (2013) and Nissenbaum (2010).

39 The rights to privacy as stated in Article 17 of the International Covenant on Civil and Political Rights is referenced in the preamble to the Privacy Act, but that right is not expressly conferred in the Australian Privacy Act or elsewhere in Australian domestic law.

The Privacy Act does not specify how a regulated entity should evaluate interests of individuals in protection of data privacy and at what level of privacy impact those interests should be adjudged to be legally protected, legitimate expectations of privacy. Sometimes it is suggested that the appropriate evaluation is whether a particular act or practice will cause a significant privacy harm to individuals. However, the Privacy Act does not state factors that a regulated entity should take into account in determining whether a particular act or practice is reasonably likely to effect a privacy harm.⁴⁰

In any event, the Act requires regulated entities to conduct a balancing of interests. Whenever the law requires balancing of interests, there is contention as to how to strike the appropriate balance. Whenever a regulated entity is required to balance its self-interest against interests of others, self-interest might be considered likely to prevail, and particularly where those others (namely, affected individuals) may not fully understand how their interests are being affected, where detection of inappropriate balancing is difficult, and where enforcement resources are stretched.

In many other jurisdictions, the domestic data privacy (data protection) statute or overarching human rights law provides a foundational legal right of privacy directly enforceable by individuals. For example, many decisions of the Court of Justice of the European Union interpreting and applying the GDPR Regulation (EU) 2016/679 (2016) commence as private litigation, often initiated by prominent privacy advocates, and turn on construction and application of Article 8 (right to respect for private and family life) of the European Convention on Human Rights. Article 8 provides:

1. Everyone has the right to respect for his private and family life, his home and his correspondence.
2. There shall be no interference by a public authority with the exercise of this right except such as is in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic wellbeing of the country, for the prevention of disorder or crime, for the protection of health or morals, or for the protection of the rights and freedoms of others.⁴¹

Without an overarching foundation or guardrail of a legal right to privacy conferred by domestic statute and enforceable by affected individuals, the Australian Privacy Act is more heavily dependent upon transparency to affected individuals as the key control or safeguard of privacy than is the case for legal rights-based privacy statutes in other jurisdictions.

40 On privacy harms, see Leonard P (June 2020) Privacy Harms: A Paper for the Office of the Australian Information Commissioner, https://www.oaic.gov.au/__data/assets/pdf_file/0012/1371/privacy-harms-paper.pdf

41 See further, the extensive case law referenced and discussed in European Court of Human Rights' Guide on Article 8 of the Convention – Right To Respect for Private and Family Life, updated on 31 August 2021, at https://www.echr.coe.int/documents/guide_art_8_eng.pdf

Entity accountability and multiparty data ecosystems

Over the last decade, data privacy reforms across the globe have rebalanced legislated privacy settings towards greater accountability of regulated entities that collect and control personal information in relation to their own acts and practices. Reform of the Australian Privacy Act can be confidently expected to follow this trend.

Many jurisdictions have recognised a distinction between 'data controllers', being entities that collect and control personal information about individuals, and data processors, being entities that process that personal information on behalf of data controllers in circumstances where the control as to subsequent uses and disclosures of that information remains with the data controller.

Those jurisdictions typically require the data controller to implement contractual safeguards and take active steps to monitor the activities of those data processors when processing personal information on their behalf. However, those jurisdictions typically do not require data controllers to actively monitor activities of entities to whom they disclose personal information where that disclosure is with the consent of the affected individual and the information then leaves the discloser's effective control.

Key requirements of the Australian Privacy Act

The Australian Privacy Act requires each regulated entity to:

- make available a privacy policy that explains generally how the entity deals with personal information about individuals
- collect personal information only as is reasonably necessary for one or more of the entity's functions or activities
- take such steps as are reasonable in the circumstances to notify affected individuals of the purposes for which the APP entity collects personal information (commonly referred to as the purpose limitation, and common across many jurisdictions)
- only use personal information for that notified purpose and related secondary purposes (commonly referred to as the secondary uses limitation, and also common across many jurisdictions), or otherwise only with informed consent of the affected individual
- collect personal information about an individual only from that individual, unless it is unreasonable or impracticable to do so, or otherwise only with informed consent of the affected individual
- obtain consent in relation to collection and uses of certain narrower categories of more 'sensitive' personal information.

The legal requirements as to 'reasonably necessary' and stated 'purpose' operate as significant constraints upon APP entities. It is therefore incorrect to characterise

the Privacy Act as principally reliant on privacy self-management by users. However, the balance between privacy self-management by users and self-responsibility and accountability of APP entities is heavily weighted towards the former.

The Australian Privacy Act does not recognise a controller–processor distinction. Regulatory guidance by the Australian Information Commissioner uses a concept of ‘effective control’ in drawing a distinction between a third-party ‘use’ of personal information at the direction of a regulated entity and provision of personal information to a third party that then is no longer acting under the direction or control of the regulated entity, being a ‘disclosure’ (Office of the Australian Information Commissioner; OAIC 2019: para 8.8, B.64).

A more recent trend in some jurisdictions has been imposition of legal accountability upon entities that curate or otherwise enable multiparty data ecosystems to monitor and control privacy-affecting activities of other entities within those multiparty data ecosystems, regardless of whether those other entities are data processors or data controllers in relation to relevant personal information. Various legal theories of responsibility and accountability of entities for acts and practices of others have been invoked, including legal theories analogous to the broad legal concept of ‘knowingly concerned’ (sanction, approve or countenance) as used in the *Australian Consumer Law*.⁴² One key issue in reform of Australian data privacy law is how to address responsibility and accountability of entities that curate or otherwise enable multiparty data ecosystems.

The illusion of (transparency and) consent

Critiques of privacy self-management mechanisms, particularly as applied to internet enabled services, focus upon:

- the impracticability of individuals reading and understanding privacy policies and requests for consent, given the volume and complexity of privacy policies and collection notices
- ‘notice and consent fatigue’, leading to users simply clicking the ‘I agree’ button without perusing or thinking about the privacy related terms. (Leonard 2020)

Many criticisms revolve around the problem of expecting affected individuals to properly understand and make a choice about whether to accept an act or practice which affects the individual’s privacy. An informed understanding requires willingness of an affected individual to engage with explanations of the *why and how* of collection, use and sharing of personal information.

42 Under the Corporations Act 2001 (section 79), Fair Work Act 2009 (section 550) and Australian Consumer Law (section 2), a person (including a company) will be ‘involved’ in a breach of the respective statutes where that person has aided, abetted, counselled or procured the contravention; or induced, whether by threats or promises or otherwise, the contravention; or been in any way, by act or omission, directly or indirectly, knowingly concerned in, or party to, the contravention; or has conspired with others to effect the contravention. See further, Australian Competition and Consumer Commission v Joystick Co Pty Ltd (2017) FCA 397 and *Yorke v Lucas* (1985) HCA 65.

Explanations provided are often technically complex. Often the counterfactual – any adverse effect on availability, quality or relevance of an internet service that an affected individual will experience if the individual does not allow data collection and uses as proposed by a service provider – is not clearly stated by the service provider. If an individual cannot understand the counterfactual, then is a clear statement as to a proposed data sharing sufficient to demonstrate individual choice to permit a relevant data flow? Individual choice requires options and informed understanding of the consequences of exercising them.

Options offered to internet users also need to be readily exercisable. If options of privacy settings are difficult to find and exercise, are they 'real' options? Some internet services offer little practical ability for a user to say 'no', or even to say 'no to that, but it might be okay if you did it this other way'.

Doubling down on consent

Some critiques suggest that the legislature should extend the categories of acts and practices for which consent is required, as well as cranking up the requirements for a valid consent. These critiques often cite with approval the EU GDPR concept of 'unambiguous express consent' (European Data Protection Board 2020). When faced with the response that such changes risk increasing the clamour for consent and resultant consent fatigue, some critics say that the impracticability of obtaining heightened consent would create disincentives for organisations from seeking consent, with an outcome of limiting privacy-affecting acts and practices of regulated entities.

Those alleged disincentives may be overstated. Jurisdictions such as Korea that had longstanding prescriptive requirements for much more granular and frequent requests for consent have not demonstrated any significant difference in privacy-affecting acts and practices of regulated entities within Korea, as compared to other, less prescriptive, jurisdictions.

The 'consent problem' under Australian data privacy law is not as acute as in other jurisdictions that have incentivised over-reliance by data controllers upon consent, which has in turn led to further erosion of the value of consent.

We should continue to contest whether and when requiring consent is sensible. We should ensure that enhancements in practical options for individuals to control their privacy settings are not compromised by any change in consent requirements. Consent should only be required, and sought, where it can be given thoughtfully, sparingly and with understanding. Consent is only 'real consent' where an individual has a real choice.

Winding back requirements regarding consent, to achieve an objective of improving data privacy, may sound both radical and counterintuitive.

Both consent fatigue and notice noise fatigue are real. Many proposals for reform of data privacy law risk doubling down on both the consent problem and the noise of policies and notices problem, casting the net too widely. We need to make consent meaningful again.

Selective noise reduction

We also need to reduce the level of 'noise' in privacy policies and privacy (collection) notices (see Leonard 2020, including references). Many privacy policies and privacy (collection) notices drown in an ocean of text, explanation of the unusual, the unexpected or the odd. For many categories of internet services, it is relatively obvious what collections and uses of personal information that are a reasonably necessary incident of provision of that service, or of offsetting the cost of provision of a no-charge or cross-subsidised service. Most consumers will understand the points-value-for-data exchange inherent in card loyalty programs, including programs offering special rewards, premium features, discounts and or privileges.

In particular, attention of consumers should be directed towards a full and fair explanation by a collector of personal information regarding sharing of that information into multiparty data ecosystems in circumstances where the entity making a disclosure statement is not in continuing control of uses and further disclosures by other entities in that data ecosystem.

Accountability of data collectors depends upon full transparency in data sharing practices. We need to ensure that each entity in multiparty data ecosystems through which personally identifying information about individuals may pass has appropriate incentives:

- to handle that information responsibly and transparently
- to not pass on information without applying appropriate controls and, in particular, to not to pass on information in a form that might reasonably be anticipated as facilitating misuse of that information by the recipient.

The right and ability of internet users to self-manage privacy settings remains important.

However, each individual should only be expected to self-manage what is realistically manageable by her or him. We should consider how to reduce the clamour of consent requests, and how to reduce the level of noise (length, technical complexity, coverage of unimportant and obvious subject matter) of privacy policies and collection notices. Noise reduction measures might include appropriately targeted exceptions, such as through legitimate interests or legitimate uses or 'compatible data practices',⁴³ or sector or application specific codes or standards, class exemptions by regulators, trust marks and certification schemes,⁴⁴ standardisation of language and use of graphics or other user-friendly transparency measures.

43 See Section 7 (Compatible Data Practice) of the Uniform Personal Data Protection Act as drafted by the US Uniform Law Commission, <https://www.uniformlaws.org/HigherLogic/System/DownloadDocumentFile.ashx?DocumentFileKey=009e3927-eafa-3851-1c02-3a05f5891947&forceDialog=0> and <https://fpf.org/blog/uniform-law-commission-finalizes-model-state-privacy-law/>

44 As in Japan, New Zealand and Singapore: Japan's PrivacyMark System is described at <https://privacymark.org>, New Zealand's Privacy Trust Mark at <https://www.privacy.org.nz/resources-2/applying-for-a-privacy-trust-mark/>, and Singapore's Data Protection Trustmark at <https://www.imda.gov.sg/programme-listing/data-protection-trustmark-certification>

Bringing it together and the role of transparency

Transparency is of course appropriate for data subjects who want to read privacy policies and collection notices. However, no regular person should be expected to read all that stuff. We need new thinking on the purposes of privacy policies and collection notices. We need less noise and clutter in our lives. Does it matter if many people don't read privacy policies and collection notices, provided that regulators, civil society organisations and potential litigants are able to do so?

However, any exception for legitimate interests, legitimate uses or 'compatible data practices' should only operate and allow a regulated entity to collect, handle or disclose personal information about individuals without consent if the processing is aligned with the ordinary expectations of affected individuals, having regard to transparent privacy policies and notices, and not harmful to direct interests of data subjects.

In particular, permitted primary purposes of collection and handling of personal information about individuals should remain subject to transparency requirements. Laws addressing fair disclosure, in terms readily understood by a reader of not unusual literacy, are a powerful deterrent against excessive or unduly intrusive data privacy practices.

Bridging the accountability gap: 'fair and reasonable' practices and organisational accountability

The effectiveness of data privacy law is questionable partly because many regulated entities have elected to adopt either a 'catch us if you can', or a 'tick the box', strategy in addressing their purported compliance with data privacy law.

Many regulated entities consider privacy risk management as another exercise in form over substance, only providing 'transparency' through buried and opaque disclosures of their privacy-affecting acts and practices.

Many data protection regulators are under-resourced, so enforcement action must be selective. Regulators have also been required to divert limited resources to address year on year increases in the number and complexity of data breaches (OAIC n.d.), and to investigating and addressing a variety of concerns about data handling practices of large online platforms and social media networks (Centre for Information Policy Leadership 2021; Future of Privacy Forum and Nymity 2018; Personal Data Protection Commission Singapore 2022:66–75).

One criticism of the Australian Privacy Act, and data privacy statutes of comparable jurisdictions, is that they do not adequately bridge the gap between ensuring:

- that there is 'transparency': a fair description is created and provided to an affected individual about the purpose and extent of a proposed data collection, use or disclosure or surveillance activity

- that this data collection, use or disclosure or surveillance activity is necessary and proportionate to achieve a reasonable outcome, with reasonableness judged by consideration of:
 - the degree of risk and extent of impact upon legitimate expectations of privacy
 - whether any individual is reasonably likely to suffer a harm that arises from this act or practice,
 - societal interests, including in health and safety of other individuals and in secure, safe and efficient operation of the internet
 - the interests of the regulated entity that wants to collect, use or disclose data and insights derived from analysis of personal information about individuals in a properly risk managed way.

Critiques often suggest that privacy self-management mechanisms need to be supplemented, or replaced, by:

- an overarching legal requirement of fairness or reasonableness (Attorney-General's Department 2021:82–93)
- demonstrated organisational accountability of the entity that is collecting, handling or disclosing personal information about an affected individual.

Differential treatment of individuals: Scoping the role for data privacy law

One major impetus for overhaul of the Australian Privacy Act and comparable statutes in other jurisdictions is increasing concern about use of personal information about individuals for differentiated treatment of those individuals.

Advances in transactor and transaction analytics, shift to online transactions, take-up of non-conventional internet-enabled devices such as personal wellness devices and smart speakers, and deployment of and rearchitecting of data platforms, have fuelled ever more sophisticated ability of service providers to use consumer data to single out an individual for differential treatment. If a supplier has reasons to single out a person – to deal, or not deal, or for a more or less favourable offer – this differentiated treatment is often possible without needing to know the identity of the person singled out.

If a supplier takes care not to know, and not to be able to work out, who it is that is being singled out, current Australian data privacy law generally doesn't regulate that singling out, or specify permissible reasons for singling out, because there is no relevant use of personally identifying information about individuals.

Differentiated treatment may be benign (positive or neutral) or have negative effects upon an affected individual. Often differentiation enables presentation of content, choices or offers that have been selected for inferred relevance or convenience. Search engines, marketplaces and comparison sites use algorithmic inferences to differentiate between users to promote presentation of particular content or choices

inferred more likely to be of interest to a user (whether or not identifiable), often with the positive effect of reducing that user's search time and effort. Regardless of operation of data privacy law, other laws limit the reasons that may motivate a supplier to single someone out for differential treatment.

An increasing variety of topic-specific and sector-specific statutory provisions regulate particular reasons for differential treatment, including laws about discrimination, consumer protection, targeting of children, tracking and surveillance, disinformation and misinformation.

One key issue for reform of the Australian Privacy Act is scoping the role for this statute in regulating profiling. Specifically, what should be regulated under this statute as a use of data in relation to an individual to single out that individual for differential treatment, and what is better addressed by the *Australian Consumer Law*, financial services laws or other topic-specific and sector-specific laws?

Even for particular applications of profiling, it may be difficult to structure the right package of new data privacy rules. A change in one area of data privacy law, such as by broadening the definition of personal information, may have substantial knock-on effects in other areas, such as increasing the complexity of technical information that needs to be disclosed, placing further stress upon consumer understanding of privacy policies and notices. Changes to settings within the Privacy Act requires consideration of the effect of a change upon the balancing of interests of regulated entities in conducting their business operations and addressing interests in privacy of affected individuals.

A key area of significant controversy is how data privacy regulation and regulators should address the most common form of algorithmically enabled differential treatment of internet users, being targeted ('programmatic' or 'personalised') digital advertising.

Targeted digital advertising as a form of profiling

Examples of digital advertising activities include using:

- information volunteered by a known (identified) consumer about their needs, preferences or interests to select and present tailored offers (for example, marketing by loyalty card program partners to card members based upon their membership data and their interactions with other program partners)
- observations of a consumer's interactions with a website (for example, searching on a travel website for flights to Cairns) to select and present offers tailored to meet a consumer's characteristics, needs, preferences or interests as inferred from those interactions (such as snorkelling gear, sunglasses and reef cruises)
- advertising services based upon entry of search terms (for example, a search for 'new kayak Sydney') to deliver advertisements to consumers searching for a related item (such as new double kayaks, life vests and paddles available in Sydney).

Digital advertising using audience segments enables 'personalisation' in the sense that a group of identified users receive digital ads targeted to address their needs, preferences or interests. However, the digital ad is not tailored to a particular recipient, and a recipient does not need to be personally identified.

Advertisers use ad networks and other adtech intermediaries to target ads to users based on characteristics such as their online behaviour, physical location, or demographics. Behavioural targeting shows ads to users based on their online activity, such as past searches or browsing history. Location-based ads target users based on where they live or when they visit a specific location, such as a stadium or shopping centre. Demographic targeting shows ads to users based on specific social categories (brackets) such as gender, income, level of activity or age.

Some adtech intermediaries also allow advertisers to target ads to custom audiences, such as previous customers. Adtech intermediaries collect data about users to create these segments, to enable these personalised ads and to measure their efficacy.

'Personalisation' is typically through creation and use of an audience segment, not individual targeting of individuals within that audience segment. For example, an adtech intermediary may offer advertisers the ability to target thousands of internet users inferred to be interested in water sports, addressable by the adtech intermediary enabling serving of ads to users that may or may not be identifiable, using technical information such as tracking codes of internet access devices and browsers.

Audience segments as used in personalised digital advertising are intended to be fit for purpose on an aggregated basis, but at the cost of some outliers: that is, overinclusion of some codes for which the 'personalisation' is not right. This 'outlier cost' often arises because an advertising services provider does not know the identity of a user, or specifics of a particular user's browsing or searching activity over time or across devices. In other words, accuracy in targeting is lost through de-identification, inferences and aggregation. However, there are privacy protective benefits of de-identification, inferences and aggregation, including:

- minimisation of collection and use of identifying details about people using internet services and browsers and devices used to interact with those services
- minimisation of sharing of data about users of internet browsers and devices: for example, an adtech services provider could offer to serve a digital ad to thousands of users of internet browsers and devices that are inferred to have an interest in outdoor water sports, without the service provider disclosing to the advertiser or the advertiser otherwise knowing the identity of these individual users or any specifics of those users' online activity.

Adtech intermediaries generally do not share personal information about individuals with advertisers. However, some operators of internet sites (publishers) collect and share personal information with adtech intermediaries, or do not monitor or control collection of personal information from their internet sites by adtech intermediaries with which they work. Some adtech intermediaries obtain personal information from

advertisers, or share personal information back with advertisers. In short, there are differing levels of compliance across the digital advertising sector with requirements and restrictions as to necessity, purpose and transparency.

Regulators around the world have expressed concerns that:

- adtech is not configured to minimise use and disclosure of personally identifying information
- internet users lack transparency, understanding and control regarding when and how their internet interactions are being tracked for the purpose of targeted advertising
- the manner of presentation and content of privacy policies, notices and requests for consent do not adequately address likely user behaviours and capabilities. (See, for example, Information Commissioner's Office 2021.)

Some publishers and digital advertising service providers have responded to these concerns. Responsive measures include improvements in clarity, simplicity and prominence of notices to internet users about ad targeting; new options for users to change tracking settings; and expansion of the subject matter categories of digital ads that they do not permit.

Other publishers and ad service providers have been slower to respond. To date, demonstrably reliable and verified implementation of good privacy practices, including privacy by design and default, have not been widely regarded as differentiators for business success, and as a result, oversharing of personal information about individuals has been common. However, this is now changing across multiple jurisdictions, through the combination of:

- consumer organisations and regulators exerting pressure upon both publishers and adtech intermediaries to adopt more privacy protective practices
- the focus of regulators broadening from acts and practices of the global digital platforms to include scrutiny of activities of other entities within the digital advertising sector
- improvements in data architectures and governance that increasingly enable less identifying information to be gathered or shared while still enabling targeted digital advertising.

In particular, adtech intermediaries are rearchitecting data handling and investing in new technologies to address oversharing of personal information, deliver more relevant ads, and prevent ad fraud. Over recent years the adtech sector has been working on transparency and accountability frameworks for sharing of attribute data across multiparty ad data ecosystems.

IAB Tech Lab is developing one example of a federated model, where each entity enabled into an ad data ecosystem would commit to transparency requirements, to observe use restrictions, to follow technical standards and assure 'privacy by default' addressable advertising and measurement.

Other proposals include substitution of tracking codes and device codes for what is variously called common ID, stable ID or universal ID. Universal ID proposals claim to provide a means by which the identity of a user, internet device or browser can be protected against being reverse-engineered to a form of identification of a user.

For example, the Prebid.js User Identity Module would enable a publisher to permit any one or more of a variety of proprietary submodule ID generators, including the TradeDesk-sponsored Unified ID (UID) 2.0, Verizon Media ConnectID, and Tapad ID, which in turn would transport or regenerate the common pseudonymous ID across other solutions. This would facilitate cookieless tracking of interactions by a unique pseudonymised user with publishers that are unrelated with each other, and also across publishers working with a variety of different adtech solution providers.

One way to address perceived intrusiveness would be to move away from creation of audience segment cohorts for targeted ads through direct correlations based upon observation of browsing behaviour of individuals.

Google sought feedback, through its Privacy Sandbox initiative, on a number of alternative technical implementations of Federated Learning of Cohorts (FLoC), whereby a user's browser is associated with a value, alongside thousands of others with a similar browsing history, which is updated over time as the collective cohort of users traverse the internet. That value, and not the user's actual browsing behaviour, is used to target ads.⁴⁵

Following feedback received by Google on FLoC, Google made substantial changes to the program and renamed it the 'Topics' API.

With Topics, an individual's browser generates inputs to machine learning algorithms that develop a cohort based on thousands of individuals' interactions, analysing URLs of the visited sites, content of pages visited and other factors. Input features to the algorithm, including the individual's browsing history, are kept local on the browser and are not uploaded.

Based on the user's activity, the browser generates rolling interest scores in a list of 300 different topic areas (for example, music or automobiles) created by Google. Those topic areas with the highest interest scores can be shared to websites through the Topics API, which can then be used for targeted advertising.

The user's Topics score is updated over time, so that it continues to have advertising utility, but (when implemented with appropriate controls) not at a frequency and without granularity of analysis that would enable direct correlations with a user's internet activity. Controls could include the ability for a publisher to opt out of inclusion in the user's list of sites for cohort calculation, individual users to opt out of inclusion within any cohort, restrictions as to uses of categories of sensitive information in creation of cohorts, and no-go zones, such as browsers used by young children.

45 Bindra C (25 January 2021) 'Building a privacy-first future for web advertising', Google Ads & Commerce Blog, <https://blog.google/products/ads-commerce/2021-01-privacy-sandbox/>

Topics implementations do not of themselves ensure responsible data governance by entities within multiparty ad data ecosystems. However, they would significantly reduce the collection and centralisation of data about an individual's internet activities, which substantially reduces availability of that data to entities within multiparty ad data ecosystems, thereby mitigating the risk of misuse of data about internet activities.

Some consumer advocates argue in favour of new legal restrictions on profiling that go well beyond existing data privacy laws. These concerns are often framed not in terms of compliance with data privacy law but in more emotive terms, such as that 'surveillance-based advertising' renders consumers 'vulnerable to manipulation, discrimination, misinformation and fraud' (for example, Norwegian Consumer Council 2021). Some of these proposals do not differentiate between segmentation of audiences for targeting of ads – that is, delivery of purely expressive content to a cohort of internet users with inferred like interests – and differentiation between users for the purpose of determining terms of dealing with an individual in relation to supply of a particular product or service.

UNSW law professor Katharine Kemp (2021) recently suggested that sharing of targeting data should be unlawful unless a consumer ticks an unticked box next to a plain message, such as 'Please obtain information about my interests, needs, behaviours and/or characteristics from the following data brokers, advertising companies and/or other third-party suppliers', with each entity named. Professor Kemp also suggested that collection should not be exempt from this rule 'simply because the companies use a pseudonym or unique identifier, rather than the consumer's given name or contact details, to link data collected by the marketplace with data about the same consumer collected by a third party'.

Such proposals would effectively preclude targeted advertising using pervasive tracking and data sharing between adtech intermediaries, whether or not using demonstrably effectively anonymisation, unless there had been an affirmative and express consent by a consumer, and then only as between entities named in that consent. Such proposals loop the debate back to the issue of consent fatigue of consumers, and whether it is reasonable to expect consumers to engage in understanding complex adtech processes.

'Purely expressive' content and 'compatible uses'

Clearly, differentiation of offers involves significant risk of unfair or illegal price discrimination, or even refusal to deal. But does the former – audience segmentation for delivery of purely expressive content to a cohort of internet users with inferred like interests – raise significant consumer protection concerns?

Reflecting this distinction, the US Uniform Law Commission's Uniform Personal Data Protection Act, published as a model law for US state legislatures, draws a distinction between compatible business practices (processing that 'is consistent with the ordinary expectations of data subjects or is likely to benefit data subjects

substantially'), which do not require a data subject's consent (but are still subject to transparency requirements) and incompatible business practices, for which either consent is required or are described 'in a reasonably clear and accessible privacy policy' as a practice 'that, unless the data subject withholds consent, will be applied by the controller or an authorized processor to personal data' (Uniform Law Commission 2022:17; Future of Privacy Forum 2021).

The model statute provides that a controller may use personal data, or disclose pseudonymised data to a third-party controller, to deliver targeted advertising and other purely expressive content to a data subject. However, a controller may not use personal data or disclose pseudonymised data to be used to offer terms, including terms relating to price or quality, to a data subject that are different from terms offered to data subjects generally: this is an incompatible data practice that requires consent, unless otherwise excepted.

Another exception addresses loyalty programs that use personal data to offer discounts or rewards:

although the targeted offering of discounts or rewards would constitute decisional treatment, these are accepted and commonly preferred practices among consumers ... This subsection does not prevent providing special considerations to members of a program if the program's terms of service specify the eligibility requirements for all participants. (Uniform Law Commission 2022)

Profiling beyond targeted advertising

Some forms of profiling-based differentiation in terms of offer involve significant risk of unfair or illegal price discrimination, or even refusal to deal. But many uses do not. An online supplier may infer the characteristics of a product or service likely to be of interest to an online user and present that user with an offering with those characteristics more quickly, or with greater prominence. 'Noise' from multiplicity of possible options is thereby decreased, with benefit to the consumer (for one expression of concerns that may arise from 'echo chambers' or other 'tunnelling', see Fish and Gal 2020).

Or a supplier may use similar data analytics capabilities to infer that an online user is less price sensitive, and elect not to offer that user as attractive a price as may be offered to other online users that are inferred to be more price sensitive.

Or a supplier may classify a user into a cohort of inferred like individuals as an exclusion audience. Consider an offer of community-rated insurance products, where an insurer has an incentive to only actively market a product to those sections of the public likely to take up the product, not less likely to make a claim under a policy.

If a health insurance product can be marketed only to an audience segment that is inferred from their recent purchases to be physically active young people

(regardless of their identity, and although that inference may be wrong in a statistically insignificant number of cases), the offer of that health insurance product may be much more profitable than if that same product is offered at that same price through broadcast media such as free-to-air television. Targeting through data inference may fundamentally alter profitability of a product or service.⁴⁶

'Minding the gap' in data privacy impact assessments

Data privacy regulation, when properly applied, should lead to a contextual assessment by each regulated entity of risks of privacy harms to individuals that may arise from acts and practices in collection and handling of data relating to persons with whom that entity deals or otherwise interacts.

A key methodology for this contextual assessment is conduct by regulated entities of data privacy impact assessment (DPIA).

Conduct of a DPIA is becoming a key feature of responsible and accountable governance and assurance of data privacy of affected individuals, including in circumstances where conduct of a DPIA is not legally mandated.

DPIAs are hard to do well, and often they are not done well. This shortcoming is increasingly problematic because DPIAs are now being repurposed as a mechanism for algorithmic or AI impact assessments, which are of necessity more complex and multifaceted than data privacy risk assessment.

Given the complexity, range and relative novelty of risks and possible mitigations that should be evaluated and addressed in a comprehensive algorithmic or AI impact assessment, it is important to ensure that DPIAs are properly adapted and applied to this new purpose.⁴⁷

Australian Privacy Principle (APP) 1 of the Australian Privacy Act 1988 requires APP entities to take reasonable steps to implement practices, procedures and systems that will ensure compliance with the APPs and enable them to deal with enquiries or complaints about privacy compliance.

In this way, the APPs require 'privacy by design', an approach whereby privacy compliance is designed into projects dealing with personal information right from the start, rather than being bolted on afterwards. Conducting a DPIA may help

⁴⁶ For an interesting analysis of possible regulatory responses to the impact of AI and advanced data analytics in the insurance sector, see Bednarz Z and Manwaring K (2021) 'Insurance, Artificial Intelligence and Big Data: Can Provisions of Chapter 7 of the Corporations Act Help Address Regulatory Challenges Brought About by New Technologies?', *Australian Journal of Corporate Law*, 36(3):216–239.

⁴⁷ To date there are relatively few published examples of fully developed tools for AI risk assessment and assurance. Many examples take the form of checklists or questionnaires rather than assurance frameworks. In December 2021, the NSW Department of Customer Service published an AI assurance framework that NSW government agencies will be required to apply from March 2022 to assess all significant projects that use bespoke AI systems before deployment. See the NSW Artificial Intelligence Assurance Framework, available at <https://www.digital.nsw.gov.au/policy/artificial-intelligence/nsw-ai-assurance-framework>

an entity to ensure privacy compliance and identify better practice. A DPIA is a systematic and documented assessment of a project that identifies the impact that the project might have on the privacy of individuals, and sets out recommendations for managing, minimising or eliminating that impact. However, conduct of a DPIA is not currently mandated by the Privacy Act.

The Privacy (Australian Government Agencies – Governance) APP Code 2017 (the Government Agencies Code) requires Australian Government agencies to conduct a DPIA for all 'high privacy risk projects' (OAIC 2020). The Government Agencies Code provides that a project may be a high privacy risk project if an agency reasonably considers that the project involves any new or changed ways of handling personal information that are 'likely to have a significant impact on the privacy of individuals'. Guidance of the Australian Information Commissioner in relation to the Government Agencies Code states:

An impact on the privacy of individuals will be 'significant' if the consequences of the impact are considerable, taking into account their nature and severity.

The consequences of a privacy impact could be significant for one individual or a group of individuals, for example, negative impacts on physical and mental wellbeing, reduced access to public services, discrimination, financial loss or identity theft. The consequences of the potential privacy impacts for a group of individuals may vary based on their individual circumstances, so you should consider whether some individuals may be more significantly impacted than others.

Sometimes projects can have a significant collective impact on society, rather than impacting on people individually. These collective impacts are likely to lead to broad public concern, for example, increased surveillance and monitoring activities, or the establishment of sensitive personal information sharing arrangements between the Commonwealth and other entities.

There is no definitive threshold to determine when an impact is 'significant' given each project will differ in nature, scope, context and purpose. Accordingly, agencies are advised to screen for factors that may raise a project's risk profile.

Environmental protection laws require entities to undertake and publish environmental impact statements addressing adverse impacts of significant development projects upon humans and the environment. Unlike the requirement to publish environmental impact statements, in most instances regulated entities are not required to publish a DPIA.

In practice, in many cases a comprehensive DPIA is not conducted, because an entity:

- makes a preliminary determination that the project does not carry significant risks of privacy harms to individuals
- determines that the Australian Privacy Act does not legally require a DPIA to be conducted
- does not recognise that it should be considering whether to conduct a DPIA.

In other cases, an entity may conduct a DPIA, but not identify and appropriately mitigate particular adverse effects on individuals as privacy harms, and accordingly leave unmitigated unacceptable residual risks of harms.

Article 35 ('Data protection impact assessment') of the GDPR covers DPIAs:

Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data.
(Regulation [EU] 2016/679 2016)

The European Data Protection Board provides the following examples of circumstances in which a DPIA should be conducted:

- if you're using new technologies
- if you're tracking people's location or behaviour
- if you're systematically monitoring a publicly accessible place on a large scale
- if you're processing personal data related to 'racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation'
- if your data processing is used to make automated decisions about people that could have legal (or similarly significant) effects
- if you're processing children's data
- if the data you're processing could result in physical harm to the data subjects.
(Data Protection Working Party 2017; European Data Protection Board 2018)

Although privacy impact assessments are becoming more common in relation to proposals for new applications and uses of personal information about individuals, there remains considerable disagreement on:

- the threshold at which a privacy impact assessment should be undertaken (that is, what is a serious risk of harm to an individual?)
- the nature and range of 'privacy harms' that should be assessed
- the criteria for assessment of risk and harm
- the level of potential risk of privacy harm and likely (or other) exposure to adverse impact at which a particular process should be assessed as requiring mitigation
- the level of residual risk of harm which is permitted to remain after appropriate mitigation.⁴⁸

48 On the relationship between DPIAs and algorithmic impact assessment, see Information Accountability Foundation (June 2021) The Road to Expansive Impact Assessments - Why It Matters at <https://b1f827.p3cdn1.secureserver.net/wp-content/uploads/2021/05/The-Road-to-Expansive-Impact-Assessments.pdf?time=1673503897>

Unlike processes for environmental assessment, the frameworks and methodologies for making a preliminary assessment of whether to conduct a DPIA, for conduct and documentation of a DPIA, and for assurance of their reliable implementation, are not yet mature. Many DPIAs are conducted as 'check the box' exercises in 'assessment-and-disclosure-washing' to ensure that disclosures match form disclosure requirements stated in privacy principles, rather than genuine attempts by entities to ensure necessity and proportionality in data handling practices, and to build privacy-by-design into those practices.

Boards and senior management often see data privacy compliance as an assurance and audit function rather than an integral and essential enabler of an entity conducting data-driven, or properly data-informed, business or other operations.

There are surprisingly few privacy professionals in some data-driven industry sectors, such as provision of digital advertising services (at least, outside of the global digital platforms and major media publishers) and adtech intermediation, and provision of digital health services. Relatively few regulated entities have developed in-house competencies of privacy professionals that are active in profit-centre lines of business.

Many in-house privacy professionals have limited opportunity to view, and participate in and influence, ongoing governance and assurance of privacy-affecting acts and practices of those entities. Privacy professionals working within entities are often sited within prudential and risk teams, rather than more directly involved in design and specification, and change management, of an entity's data architectures and data handling practices. As a result, significant privacy-affecting practices can creep in unassessed into an entity's ways of working and dealing with individuals, even within entities that otherwise properly conduct privacy impact assessment upon initiation of new major projects.

Often privacy risk assessment is:

- episodic, conducted only upon initiation of major new projects
- outsourced to the fast-growing information risk practices of the big consultancies, with one result being that privacy risk assessment is often subsumed within, and obscured by, primary focus upon information security risk assessment (outsourcing may also lead to an entity failing to develop in-house competencies in risk of harms assessment or to embed those competencies in its business-as-usual processes).

The most common failing of DPIAs is that they are point-of-time and often not revisited and revised as a project or product development progresses and pivots, or to take into account how a product or service is deployed and used over time.

As agile methodologies for design and development become more commonly used, and product and service life cycles shorten, the likelihood of misfit between a DPIA and reality increases. For example, identifiability risk changes unpredictably over time. Individual level transaction and transactor datasets relating to humans at any

particular point of time and within a particular data environment sit at a point within a spectrum (continuum) of identifiability from identified, to reasonably identifying (that is, pseudonymised), to effectively or functionally anonymised, to pervasively anonymised.

Information may shift, or be shifted, towards ends of the spectrum, depending on factors including:

- specifics of the processing – for example, sensitivity of the variables in the original dataset, techniques used to reduce the identifiability of individuals in the data, and analytical methods or processes used (that is, use of pattern matching to single out unique transactors)
- the data environments involved – for example, the technical and organisational measures put in place to control access to the data and reduce identifiability risk
- an entity's risk management process – for example, how an entity identifies and mitigates re-identification risks in the processing.

A comprehensive DPIA conducted at a point of time should assess identifiability risk having regard to both the nature of the data and the environment in which that data is held and processed.

Often focus upon identifiability of data on the face of the data itself distracts attention of decision-makers from the specification of controls and safeguards to be applied over the environments in which that data is held and processed. Many DPIAs do not appropriately address and specify the environment in which that data is held and processed, nor do they lead to implementation of change control within an entity to detect and appropriately address any significant change in the environment in which that data is held and processed. Environmental factors include:

- additional data that may exist (for example, other databases, personal knowledge, publicly available sources)
- who is involved in the processing, and how they interact
- the operational governance processes that are in place to control how the information is managed (for example, who has access to it, for what purposes, and whether unauthorised accesses or uses will be promptly detected)
- contractual and other legal considerations that may apply, such as effective gateways that may impact the potential for disclosing information that enables individuals to be identifiable, and prohibitions that have the effect that while information could technically be combined to aid identifiability, doing so is against the law (for example, professional confidentiality). (Elliot et al. 2020)

Over time, operation of various factors may cause information to shift towards ends of the identifiability spectrum. For example:

- new information may be brought into the data analytics environment, increasing susceptibility to mosaic or pattern re-identification

- new external information may become reasonably available to a person attempting to re-identify individual data, also increasing susceptibility to mosaic or pattern re-identification
- new threat vectors may emerge
- new technological means to re-identify an individual may become available to threat vectors
- verification of operation of technical and operational controls and safeguards may break down, or levels of training or adherence to operational controls may decline, so that processes and practices become more risky.

DPIAs should be a valuable tool for regulated entities to ensure that their acts and practices in handling data relating to consumers and other citizens do not create significant risks of privacy harms to individuals. Too often, the tool is not used, or is used poorly, or is not brought out again when the tool needs to be used again.

General challenges and guiding principles for the responsible adoption of automated decision-making

Reforms of the Privacy Act will not address many concerns about harms to individuals, or to society, potentially arising from applications of new technologies and advanced data analytics. Artificial intelligence (AI), machine learning (ML) and other algorithmic inference engines, and collection of non-traditional data (for example, through IoT devices and other smart cities and smart infrastructure applications) also give rise to concerns that should be addressed by responsible innovators.

Concerns include:

- legal and regulatory compliance
- competent use and adequate human oversight
- an entity's ability to explain decisions made with AI or other algorithmic automation systems to the individuals affected by them
- reduction in an entity's ability to be responsive to customer requests for information, assistance, or rectification
- social and economic impacts.

Some concerns are specific to advanced inference engines such as ML. Others arise from more basic algorithmically driven differentiation between users/consumers/citizens. For example:

- The performance of AI systems and other algorithmic inference engines crucially depends on the quality of the data used. However, data quality issues can be difficult to identify and address.
- Models developed with ML can have characteristics that set them apart from more conventional models, including opaqueness, non-intuitiveness, and adaptivity.

- Adoption of AI and automated decision-making by organisations is often accompanied by significant changes in decision-making processes within organisations, creating risks of over-reliance (dependency upon AI in making decisions in contexts or scenarios where that AI is not reliable) and opacity as to why decisions are made.

Adoption of AI and automated decision-making can be accompanied by significant changes in the structure of technology supply chains, including increases in supply chain complexity and the reliance on third-party providers. Focus upon AI outputs risks creating a frame of review that underestimates or ignores how humans using AI may rely upon AI outputs to effect outcomes that are not fair, socially responsible, reasonable, ethical or legal.

The use of AI and automated decision-making can be accompanied by an increased scale of impacts when compared to conventional ways of performing business tasks. When things go wrong, unintended consequences can be very significant and very rapid.

Recent years have seen a rapidly growing literature on AI ethics principles to guide the responsible adoption of AI and automated decision-making, variously described but often reduced to fairness, accountability, transparency, equity and safety/sustainability (FATES or FEATS).⁴⁹

As noted by Dr Florian Ostmann and Dr Cosmina Dorobantu of the Alan Turing Institute (Ostmann and Dorobantu 2021), the general challenges that AI poses for responsible innovation, combined with the concrete harms that its use in financial services can cause, make it necessary to ensure and to demonstrate that AI systems are trustworthy and used responsibly.

AI transparency – making information about AI and automated decision-making systems available to relevant stakeholders – is fundamental to both of these needs. Transparency acts as an essential precondition, an enabler, for ensuring that other principles for responsible AI are met. Transparency is therefore a logical first step for explainability (Information Commissioner's Office and The Alan Turing Institute (n.d.); The Alan Turing Institute 2019) and for responsible and accountable deployment of AI and other automated decision-making systems. Governance and assurance frameworks and processes, and feedback and reassessment loops, depend upon transparency.

Information about AI and automated decision-making systems can take different forms and serve different purposes. A holistic approach to AI transparency involves giving due consideration to different types of information, different types

49 The AI Ethics Guidelines Global Inventory, a project by AlgorithmWatch, maps frameworks that seek to set out principles of how systems for automated decision-making (ADM) can be developed and implemented ethically. The database currently includes 173 guidelines: <https://inventory.algorithmwatch.org/>. See further, Australian Computer Society (2021) *The Ethics And Risks Of AI Decision-Making*, and the reports and other resources listed there (under 'Further reading' at pages 26 to 29), at <https://www.acs.org.au/insightsandpublications/reports-publications/the-ethics-and-risks-of-ai-decision-making.html>

of stakeholders, and different reasons for stakeholders' interest in information.

Transparency needs include access to reliable information about:

- a system's logic (system transparency)
- the processes surrounding a system's design, development and deployment (process transparency)
- how a system is used and relied upon as a component in a decision-making chain and in different contexts and scenarios for decision-making (contextual decision transparency)

by:

- personnel in different roles within the organisation using the system (internal operations and oversight transparency)
- external stakeholders such as regulators (external oversight transparency)
- external stakeholders, such as citizens on whom use of AI or other automated decision-making may cause significant effects, and civil society organisations and regulators, to enable those stakeholders to understand possible adverse effects such as overly granular profiling or unfair differentiation between individuals, or excessive surveillance (external affected individuals' transparency).

For system and process transparency alike, there are important questions about how information can be obtained, managed and communicated in ways that are intelligible and meaningful to different types of stakeholders.

Both types of transparency – internal and external – are relevant in ensuring and demonstrating that applicable concerns are addressed effectively. These concerns may arise regardless of whether the application of AI involves any use of personally identifying information about (identified or identifiable) users, or information about pseudonymised user-specific activities or behaviours; and regardless of whether creation and use of the cohort involves unlawful discrimination or other infringement upon currently legally recognised human rights.

In other words, compliance with existing data privacy and anti-discrimination laws is a relevant concern, but only one concern.

Application of broader principles of fairness, equity, accountability and transparency in uses and applications of data about individuals – and not just personal data about these individuals – must become an essential feature of processes and practices of data governance and assurance of businesses, government agencies, political parties and not-for-profits.

Changes in the scope and coverage of the Australian Privacy Act are necessary to address some of these concerns. However, data privacy law is not the right instrument to address many concerns about harms to individuals, or to society, potentially arising from applications of new technologies and advanced data analytics.

In addition, until we better articulate those concerns, and good practice to address them, we cannot fully assess whether other new laws are necessary to cover the broad range of entities that are now deploying AI, ML and other automated decision-making, and collecting non-traditional data.

Before we condemn entities for failing to be ethical or socially responsible, or impose broad regulatory constraints across diverse applications, we need to ensure that entities applying these new technologies and data analytics capabilities understand how they can ensure that they reliably and verifiably evaluate what they should, or should not, be doing. Publication of ethical principles, without more, is simply not good enough: we need to provide clear guiderails for regulated entities. Principles of ethical, or socially responsible, conduct will not be consistently and reliably translated into practice unless there is also clear articulation of:

- what good practice looks like
- how good practice should be assessed and given effect through methodologies and tools
- how unacceptable or illegal practices will be detected and prevented
- how to achieve the right balance between incentives for good behaviour and sanctions for unacceptable behaviour.

Challenging indeed, but this is a necessary concomitant of continuing to derive manifest benefits to society from applications of new technologies and advanced data analytics.

Concluding remarks: a design manifesto for an Australian Privacy Act that is fit for purpose in the 21st century

The following contentions might inform redesign of the Australian Privacy Act:

1. Federal, state and territory data privacy statutes in Australia, and in many other jurisdictions, are no longer fit for purpose. Regulation focuses upon ensuring that regulated entities provide transparency to individuals and afford those individuals with (alleged) choice regarding collection, uses and disclosure of personally identifying information about them. Choice is often illusionary. Regulation does address reasonable necessity to effect a stated purpose but does not squarely address reasonableness or proportionality of acts and practices of regulated entities in collecting, handling and disclosing personal information about individuals.
2. 'Because of the Privacy Act' is enabled as an excuse to impede individual-level data linkage for population analytics conducted for societal benefit, even if conducted with appropriately isolated and controlled and safeguarded data analytics environments.
3. Risk of privacy harms to individuals should not be discounted, but societal benefit also needs to be accorded due weight. Many claims of social

beneficence and appropriate controls by would-be data analytics entities do not pass objective assessment. However, other data analytics projects that implement best practice governance and assurance are impeded, delayed and often rendered impracticable. Ethics review and approval processes are cumbersome, episodic (project orientated, not enabling standing up of continuing controlled data environments) and not sufficiently informed by preceding approval conditions. Too many ethics committees spend too much time in well intentioned reinvention of the wheel, which could be avoided if conditions devised for prior analogous reviews were readily available to inform the committee's deliberations.

4. Federal, state and territory statutes addressing use of surveillance and tracking devices are difficult to interpret and apply in relation to emerging technologies and novel uses of geolocation data, biometrics and pattern analysis to differentiate between persons in how they are dealt with. Some provisions in those statutes are inconsistent, with the effect that it is often impracticable to deploy uniformly across Australia a service that uses surveillance or tracking technology.
5. The interaction between federal, state and territory data privacy statutes, health information statutes, and surveillance and tracking devices statutes, is increasingly problematic. Health-related data is tied up in a labyrinthine interaction of federal, state and territory regulation and regulators. Many innovative health and IoT (smart utilities and smart infrastructure) applications require interactions of these regulatory schemes to be addressed with multiple agencies, for no manifest benefit in assessment and mitigation of risk of harms to affected individuals or protection of the public interest. For data uses and sharing, there is no 'one-stop (regulatory or regulator) shop', and very limited mutual recognition, across Australian jurisdictions.
6. Because data privacy statutes are focused upon acts and practices in handling of information in relation to reasonably identifiable individuals, these statutes generally do not address other data and surveillance applications that enable entities to differentiate on their treatment of persons based upon observations or inferences made by those entities as to characteristics, behaviours, interests or attributes of individuals, or small cohorts of individuals, that are not reasonably identifiable.
7. Emerging technologies increasingly enable collection and use of data that facilitates real time and granular differentiation between non-identified individuals, or grouped cohorts of 'like individuals', to enable entities using those technologies to work out whether, how or on what terms to deal with individuals. If individuals are not reasonably identifiable, this collection and use of data is not an act or practice currently regulated under data privacy statutes as a handling of personal information.
8. In many situations non-identifying differential treatment of persons is benign. Differential treatment of unidentifiable persons generally does not cause

significant risk of privacy harms to affected individuals and should not be regulated, because entities should be incentivised to ensure that persons remain unidentifiable. In some situations, this differentiation is beneficial to an affected individual, by enabling more efficient provision of content, or an offer or delivery of products or services.

9. In any event, consumer protection and anti-discrimination laws address many forms of differentiation between consumers rightly considered unfair or otherwise illegal. Privacy regulation should remain focused upon risk to individuals of privacy harms. It should not displace appropriate development of broad form consumer protection law, or the making (where there is good policy justification) of topic and sector specific statutory provisions to regulate other non-identifying differential treatment, such as laws addressing particular forms of unlawful discrimination, targeting of children for unhealthy or otherwise inappropriate content or products, excessive surveillance, disinformation and misinformation.
10. In considering reform of Australian data privacy statutes, we need to go back to basics and ask 'what harms should privacy law address?', or as Professor Julie Cohen put it, 'what privacy is for' (Cohen 2013). Revised data privacy statutes should afford due weight to ensuring that Australian society derives benefits from applications of advanced data analytics and AI, and from socially beneficial data sharing, while also ensuring that regulated entities are accountable for mitigating risk of privacy harms to individual humans and enabling humans to go about their lives without excessive intrusion upon reasonable expectations of seclusion.
11. Protection of data privacy interests of individuals requires an approach that combines top down (what is privacy?) and bottom up (what harms are we seeking to avoid or mitigate and manage?). This conclusion does not lead us to a crisp definition of data privacy. Alas, the search for crisp statutory definitions of privacy, and privacy harm, is a search for a chimera. This conclusion explains why almost all data privacy statutes refer to a right of individuals in and to (data) privacy, and to be protected against (data) privacy harms, without telling us much more about what privacy and a privacy harm actually mean.
12. The foundation of most modern data privacy statutes – notice to affected individuals and affirmative consent as to more privacy-affecting activities – remains relevant. However, we need new clarity of thinking on the purposes of privacy policies and privacy (collection) notices, to reduce the information burden upon affected individuals. We all need less clutter in our lives. Most paragraphs in most privacy disclosures are unnecessary noise. Whether it is through legitimate interests, industry standards, class exemptions by regulators, or brave new concepts such as compatible data practices, we need to reduce the level of noise in privacy policies and notices.
13. Citizens should only be expected to self-manage what is realistically manageable by them. Current regulation encourages erosion of the value of

consent. Many proposals for reform of data privacy law risk doubling down on the problem, casting the net of consent too widely. Consent should only be sought where it is reasonable to believe it will be given (or withheld) actively, thoughtfully, sparingly and with understanding.

14. Any exception for legitimate interests, legitimate uses or 'compatible data practices' should only operate and allow a regulated entity to collect, handle or disclose personal information about individuals without consent if the processing is aligned with the ordinary expectations of affected individuals, having regard to transparent privacy policies and notices, and not harmful to direct interests of data subjects. In particular, permitted primary purposes of collection and handling of personal information about individuals should remain subject to transparency requirements.
15. Some proposals for reform of data privacy laws respond to shortcomings of the notice and consent framework by advocating new measures of organisational accountability, including objective fairness or reasonableness of data privacy practices. There is an important role for organisational accountability in data privacy law.
16. One key issue in reform of Australian data privacy law is how to address responsibility and accountability of entities that curate or otherwise enable multiparty data ecosystems that share information about activities and attributes of citizens. Addressing this concern requires measures that combine increased transparency to affected individuals with organisational accountability. Introduction of a 'data controller–data processor' distinction into the Australian Privacy Act might assist in reducing clutter and noise in privacy disclosures and improve understanding of regulated entities regarding their responsibilities in management and oversight of data ecosystems that those entities enable or operate.
17. Attention of consumers should be directed towards full and fair explanation by a collector of personal information as to the sharing of that information into multiparty data ecosystems, particularly in circumstances where the entity making a disclosure statement is not in continuing control of uses and further disclosures by other entities in that data ecosystem.
18. The right and interests of individual humans to go about their lives without excessive intrusion upon reasonable expectations of seclusion needs the protection of a data privacy regulator that is credibly resourced, empowered and focused. We should be realistic and ensure that regulated entities have appropriate incentives to be responsible in, and accountable for, their acts and practices in handling of personal information. Regulatory incentives include real likelihood that an empowered and resourced data privacy regulator will take enforcement action and seek sanctions.
19. The data privacy regulator should also be empowered and resourced to issue detailed guidance and to consult with regulated entities about good data privacy governance, privacy protective processes and data assurance practices.

20. Protection of consumers from unfair contact terms and deceptive trading practices requires a consumer protection regulator of like attributes and qualities. There is significant overlap. Continuing discussions on alignment between these regulators will be necessary, but they fulfil different functions. When data privacy is seen as a consumer protection function, we have forgotten what data privacy is for.

References

- The Alan Turing Institute (2 December 2019) Project ExplAlin. <https://www.turing.ac.uk/news/project-explain>
- Attorney-General's Department (October 2021) Privacy Act Review Discussion Paper. <https://consultations.ag.gov.au/rights-and-protections/privacy-act-review-discussion-paper/>
- Centre for Information Policy Leadership (July 2021) How the "Legitimate Interests" Ground for Processing Enables Responsible Data Use and Innovation. https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/cipl_white_paper_-_how_the_legitimate_interests_ground_for_processing_enables_responsible_data_use_and_innovation__1_july_2021_.pdf
- Cohen JE (2013) 'What Privacy is For', *Harvard Law Review*, 126:1904–1933.
- Data Protection Working Party (13 October 2017) Guidelines on Data Protection Impact Assessment and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679 (wp248rev.01), European Parliament and of the Council. https://edpb.europa.eu/our-work-tools/our-documents/guideline/data-protection-impact-assessments-high-risk-processing_en
- Elliot M, Mackey E and O'Hara K (2020) *The Anonymisation Decision-Making Framework: European Practitioners' Guide*, 2nd edition, UKAN, University of Manchester.
- European Data Protection Board (25 May 2018) Data Protection impact assessments High risk processing. https://edpb.europa.eu/our-work-tools/our-documents/guideline/data-protection-impact-assessments-high-risk-processing_en
- European Data Protection Board (2020) Guidelines 05/2020 on consent under Regulation 2016/679. https://edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-052020-consent-under-regulation-2016679_en
- Fish E and Gal M (16 March 2020) 'Echo Chambers and Competition Law: Should Algorithmic Choices be Respected?' [advanced online publication of chapter from Charbit N and Ahmad S (eds) Frédéric Jenny Liber Amicorum: Standing Up for Convergence and Relevance in Antitrust Volume II]. <https://ssrn.com/abstract=3555124>
- Future of Privacy Forum (21 July 2021) Uniform Law Commission Finalizes Model State Privacy Law. <https://fpf.org/blog/uniform-law-commission-finalizes-model-state-privacy-law/>
- Future of Privacy Forum and Nymity (2018) Processing Personal Data on the Basis of Legitimate Interests under the GDPR: Practical Cases. https://fpf.org/wp-content/uploads/2018/04/20180413-Legitimate-Interest_FPF_Nymity-2018.pdf
- House of Commons and House of Lords Joint Committee on Human Rights (3 November 2019) *The Right to Privacy (Article 8) and the Digital Revolution*, HC 122 HL Paper 14, UK Parliament. <https://publications.parliament.uk/pa/jt201919/jtselect/jtrights/122/122.pdf>

Information Commissioner's Office and The Alan Turing Institute (n.d.) Explaining decisions made with AI. <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-ai/>

Information Commissioner's Office (25 November 2021) Information Commissioner's Opinion: Data Protection and Privacy Expectations for Online Advertising Proposals. <https://ico.org.uk/media/about-the-ico/documents/4019050/opinion-on-data-protection-and-privacy-expectations-for-online-advertising-proposals.pdf>

Kemp K (17 August 2021) 'How One Simple Rule Change Could Curb Online Retailers' Snooping on You', The Conversation. <https://theconversation.com/how-one-simple-rule-change-could-curb-online-retailers-snooping-on-you-166174>

Leonard P (June 2020) Notice, Consent and Accountability: Addressing the Balance Between Privacy Self-Management and Organisational Accountability: A Paper for the Office of the Australian Information Commissioner, Office of the Australian Information Commissioner. https://www.oaic.gov.au/__data/assets/pdf_file/0003/2010/notice-and-consent-paper-for-oaic.pdf

Nissenbaum H (2010) Privacy in Context: Technology, Policy, and the Integrity of Social Life, Stanford Law Books, 2010.

Norwegian Consumer Council (June 2021) Time to Ban Surveillance-Based Advertising: The Case Against Commercial Surveillance Online. <https://www.forbrukerradet.no/wp-content/uploads/2021/06/20210622-final-report-time-to-ban-surveillance-based-advertising.pdf>

Office of the Australian Information Commissioner (n.d.) Notifiable data breaches statistics. <https://www.oaic.gov.au/privacy/notifiable-data-breaches/notifiable-data-breaches-statistics>

Office of the Australian Information Commissioner (July 2019) Australian Privacy Principles guidelines. <https://www.oaic.gov.au/privacy/australian-privacy-principles-guidelines>

Office of the Australian Information Commissioner (14 September 2020) When do agencies need to conduct a privacy impact assessment? <https://www.oaic.gov.au/privacy/guidance-and-advice/when-do-agencies-need-to-conduct-a-privacy-impact-assessment>

Ostmann F and Dorobantu C (11 June 2021) AI in Financial Services, The Alan Turing Institute. 2021, <https://doi.org/10.5281/zenodo.4916041>

Personal Data Protection Commission Singapore (16 May 2022) Advisory Guidelines on Key Concepts in the Personal Data Protection Act. <https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Advisory-Guidelines/AG-on-Key-Concepts/Advisory-Guidelines-on-Key-Concepts-in-the-PDPA-17-May-2022.pdf>

Regulation (EU) 2016/679 (27 April 2016) On the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), European Parliament and Council.

Uniform Law Commission (2022) Uniform Personal Data Protection Act, Final Act with comments. <https://www.uniformlaws.org/HigherLogic/System/DownloadDocumentFile.ashx?DocumentFileKey=009e3927-eafa-3851-1c02-3a05f5891947>

United Nations High Commissioner for Human Rights (13 September 2021) The right to privacy in the digital age. <https://www.ohchr.org/en/documents/thematic-reports/ahrc4831-right-privacy-digital-age-report-united-nations-high>

Chapter 6

Final Words



Pleasurable leap

An older, shining child sings beyond the tiger

It is a truism that technology generally moves faster than public opinion, and both move faster than regulatory change.

This collection of views is just that, five voices and five perspectives on the future world of data use, implications for privacy and consent in a rapidly rising digital tide. A common perspective in these five views is the need to address the true nature of these challenges before considering regulatory reforms.

Currently scoped data privacy laws and consumer protection laws are not the appropriate frameworks to address some of key challenges of new applications of data sharing, advanced data analytics and AI/ML affecting humans and the environment. Socially beneficial applications need to be accommodated, without creating workarounds of legal protections of consumer rights and expectations of data privacy. As the EU has recognised in proposals for new regulation of AI, addressing adverse impacts upon some groups of citizens of differentiated treatment of citizens enabled through algorithmic individuated effects requires fresh policy thinking and new regulation.

Fiddling at the edges is not what is required, rather addressing the fundamental issues of what fundamental concepts such as privacy and consent mean in an on-line, hyperconnected, data prolific world.

Ultimately it is about building trustworthy frameworks for data use. As the world continues to deal with the ongoing challenges of a global pandemic and climate change, the erosion of public trust accompanying these crises appears to be accelerating calls for fresh thinking about ways to work with and for our communities. In the ecosystem of data sharing and use, for instance, building and maintaining public trust is essential for maintaining public confidence in the way that data (especially public data) is being used. Alongside this "awakening" to the value of data when shared, we see rising concerns about the governance around data sharing, especially how to avoid making mistakes or taking unnecessary risks with personal and private data.

While the world increasingly awakens to the new post COVID "normal" in social and work practices, one thing is certain: change can come very quickly and from unexpected sources. Future shocks will occur and accelerate trends which shape our world in the longer term. Technology, digitisation and AI will continue to disrupt the industries we serve. Our changing population profile will shape the world for decades to come as will changing climate. The next global pandemic or next global shock may mean that once again, the need to adapt to change is accelerated from years to weeks.

Thanks to Poem Generator for creation of Haikus used in this publication. Thanks also to NightCafe, an Australian AI art generator that made the cover images for this book.⁵⁰

50 <https://www.poem-generator.org.uk/haiku/> and <https://creator.nightcafe.studio/>





About the Australian Computer Society

ACS is the professional association for Australia's technology sector.

We represent technology professionals across industry, government and education. Our aim is to grow the nation's digital skills and capacity.

Wherever you may be in your tech career, ACS has the solution to suit your needs and take your career forward.

Plan your career

Assess and profile your current skills, understand your competencies, get recognised as a Certified Professional and map your career plan.

Learn new skills online

Gain new skills across cyber security, cloud tech, AI, machine learning and more, with over 8,000 flexible online videos and courses.

Grow your tech network

Meet the right people – network with other tech professionals as well as leaders from some of the biggest local and global organisations.

Stay up to date and relevant

Stay informed on industry trends and emerging technologies with over 200 events, masterclasses, research projects and case studies.

Be inspired by industry leaders

Join mentoring programs designed to accelerate your career growth. ACS mentors are leaders who are here to help guide you.

Protect yourself

Stay protected with comprehensive liability insurance.

Have a voice

On behalf of tech professionals ACS engages with media and policy makers on the issues affecting the technology sector, along with providing a range of resources to educators and industry to boost the nation's digital capabilities and competitiveness.

**Unlock your potential – find out more
about joining ACS at [acs.org.au](https://www.acs.org.au).**

Contact us

General enquiries

E: info@acs.org.au

T: +61 2 9299 3666

