Submission of

Shaun Chung and Rohan Shukla

To the Privacy Act Review

November 2020

1. Introduction and summary

We're inventing things much faster than we're considering the legal, societal, and ethical implications. The implications of the digital economy are vast and complex and, as it continues its rapid growth in unknown ways, many more challenges lie ahead. The *Digital Platforms Inquiry* and this *Privacy Act Review* (**Review**) are being conducted at an inflection point in our society, where privacy has become one of the defining subjects of the 21st century. We commend those involved; tackling privacy challenges in an ever-evolving landscape is no easy feat.

But we have to stop playing catch-up. Today, the digital economy is already so prevalent that we can no longer think of digital and physical as separate. By the time the Review's completed and the Privacy Act is upgraded for the digital economy, we'll probably be staring down the barrel of myriad new challenges from next-generation developments.

Claude Shannon, the father of information theory, famously said that "we may have knowledge of the past but cannot control it; we may control the future but have no knowledge of it." We need to start taking control of privacy protections in a manner that remains resilient in an unknown future. To do that, we have to start thinking about privacy with a long horizon, and break out of the cycle where Australian privacy laws continue to chase technological advancements and societal shifts. If we fail to seize this opportunity, it may be impossible to put the proverbial genie back in the bottle.

This submission addresses the following key items:

- **Privacy is a fundamental right that must be protected.** A majority of Australians, like many people globally, are increasingly concerned about privacy and have experienced privacy breaches. Privacy must be protected to the highest possible degree, and the central privacy legislation in Australia shouldn't merely *promote* the protection of privacy nor balance the interests of entities with privacy.
- **Privacy protections should be expanded.** 'Personal information' should be broadened to include all information related to an individual or an identifiable individual (which is binary and shouldn't be subject to reasonableness). Inferred personal information is private and should be expressly protected, too.
- **We need better accountability and transparency with respect to inferred personal information.** Predictive technology has enabled entities to make calculated guesses with a high degree of confidence. These guesses drive decisions that have real-world consequences ranging from mild inconveniences to, at the far end of the spectrum, severe adverse impacts *at scale*. So, entities shouldn't be allowed to discover and use inferred personal information in obscurity without parameters. More discipline's required and can be implemented by, as a first step, extending privacy protections under the Privacy Act to inferred personal information.
- **Privacy protections should also extend to certain types of non-personal information.** Today, it's easy to re-identify individuals to which anonymised information relates and there are strong incentives to do so. Anonymised information no longer protects privacy. Therefore, various privacy protections should extend to non-personal information used in re-identification or is otherwise sensitive. We should also shine a light on the shadow industry of data brokerage.
- **We need better ways to collect informed consent.** Visual privacy summaries could supplement long-form privacy notices, using standardised privacy icons designed by the OAIC in collaboration with interested persons. These summaries will also help inform the large (and growing) number of internet users who are under 18.

We look forward to continuing the conversation on protecting privacy in Australia, and considering other submissions and the Australian Government's findings from this Review.

2. Objectives of the Privacy Act

1. Should the objects outlined in section 2A of the Act be changed? If so, what changes should be made and why?

Yes. The objects set out the purpose, general aims, and principles of the Privacy Act.[1] In the 2020 ACAP survey, over 70% of respondents see privacy as a major concern, and 83% of respondents want tougher privacy laws. To an extent, this means the objects of the Privacy Act have proven insufficient and don't reflect the gravity of the subject matter. Therefore, we need to rethink them.

Protection of privacy shouldn't merely be promoted; privacy is a fundamental right that must be protected.

Section	Current	Proposed
2A(a)	to promote the protection of the privacy of individuals	to protect the fundamental right of natural persons to the protection of personal information[2]
2A(d)	to promote responsible and transparent handling of personal information by entities	to ensure responsible and transparent handling of personal information by entities

The Privacy Act should recognise a fundamental right to privacy. The *Universal Declaration of Human Rights, the International Covenant on Civil and Political Rights* (**ICCPR**), to which Australia is a signatory, recognises a human right to privacy. Additionally, the constitutions of over 185 countries around the world contemplate a right to privacy.[3] Of course, Australia is not compelled to codify a right to privacy by virtue of being an ICCPR signatory and each of those 185 countries will have varying standards for protecting privacy. But it demonstrates the global sentiment that privacy is, or has become, a fundamental right that needs to be protected. We believe the Australian Government is well-aware of a similar sentiment in Australia,[4] so it isn't necessary to dive into it here.

The purpose of the Privacy Act should be to protect privacy, not just promote it. 'Promoting' gives the impression that the Privacy Act merely guides and encourages, and entities are trusted to do the right thing.

[1] ALRC, 'The objects of the Act', *For Your Information: Australian Privacy Law and Practice (ALRC Report 108)* (Webpage, 16 August 2010) <https://www.alrc.gov.au/publication/for-your-information-australian-privacy-law-and-practice-alrc-report-108/5-the-privacy-act-name-structure-and-objects/the-objects-of-the-act/>. The ALRC explored the purpose of objects in legislation and recommended that they are included in a previous version of the Privacy Act.
[2] The proposed object is borrowed from the GDPR, Article 1, paragraph 2, which states: "This Regulation protects fundamental rights and freedoms of natural persons and in particular their right to the protection of personal data."
[3] See Constitute Project, 'Right to Privacy', *Constitute Project* (Webpage) <https://www.constituteproject.org/search?lang=en&key=privacy>. The Constitute Project was developed by the Comparative Constitutions Project founded by constitutional scholars in the United States to centralise the world's constitutions.
[4] ALRC, *Report into the Serious Invasions of Privacy in the Digital Era* (2014). This report explored recognising a right to privacy at law and balancing that right with freedom of expression and journalistic freedoms, among other things.

- **This standard simply isn't high enough.** Privacy is a fundamental and essential value that must be protected.
- **Leaving too much to the entities' discretion creates uncertainty.** Let's look at Apple, Amazon, Facebook, Google, Microsoft, and Netflix which apparently collectively account for 43% of internet consumption.[5] Of the six, two (Apple and Microsoft)[6] expressly recognise a fundamental right to privacy, which is a proxy for a certain standard for protecting privacy. But the remaining four (i.e. Amazon, Facebook, Google, and Netflix), like many other organisations, don't expressly recognise a fundamental right to privacy. Given the seemingly subtle differences, we'll need to untangle the nuances of each organisation's privacy posture to assess how they practically affect us. This is burdensome and it's impractical to expect Australians to conduct this assessment for every entity they interact with.

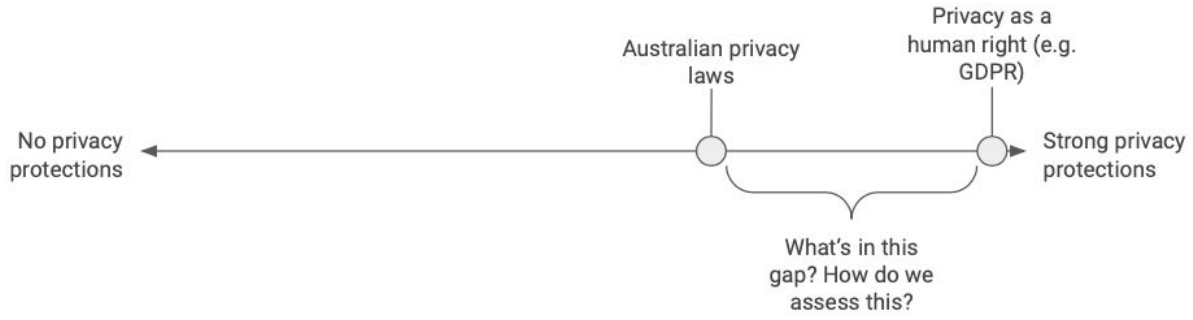

- **We can't just take entities at their word.** Even the most well-meaning entities can't be absolutely sure that their partners treat privacy in the same way or that personal information they've collected won't fall into the wrong hands. Also, there are those who will push the boundaries in pursuit of self-interest — over 50% of respondents in the 2020 ACAP survey experienced a problem with how their data was used in the 12 months to April 2020. Unsurprisingly, an estimated 93% of Australians wouldn't trust organisations (presumably non-governmental) to protect their privacy.[7]

Protecting privacy, and interests of entities, is not a balancing act.

Section	Current	Proposed
2A(b)	to recognise that the protection of the privacy of individuals is balanced with the interests of entities in carrying out their functions or activities	to ensure the protection of the privacy of individuals by entities, having regard to the interests of entities in carrying out their functions or activities

[5] See Sandvine Incorporated, 'Over 43% of the internet is consumed by Netflix, Google, Amazon, Facebook, Microsoft, and Apple: Global Internet Phenomena Spotlight' (Webpage, 30 August 2019) <https://www.sandvine.com/blog/netflix-vs.-google-vs.-amazon-vs.-facebook-vs.-microsoft-vs.-apple-traffic-share-of-internet-brands-global-internet-phenomena-spotlight>.
[6] See International Association of Privacy Professionals (**IAPP**), 'Big Tech's Shift to Privacy' (Webpage) <https://iapp.org/resources/ article/big-techs-shift-to-privacy-2/>. See also Apple, 'Privacy - Overview' (Webpage) <https://www.apple.com/au/privacy/>, which states that "Privacy is a fundamental human right. At Apple, it's also one of our core values." See also Microsoft, 'Microsoft's GDPR Commitments to Customers of our Generally Available Enterprise Software Products' (Webpage) <https://docs.microsoft.com/en-au/legal/gdpr>, which states that "At Microsoft, we believe privacy is a fundamental right".
[7] See ITWire, 'Australian consumers put a price on privacy: almost half would pay more to do business with an organisation committed to protecting their personal data' (Webpage) <https://www.itwire.com/guest-articles/australian-consumers-put-a-price-on-privacy-almost-half-would-pay-more-to-do-business-with-an-organisation-committed-to-protecting-their-personal-data.html>. OpenText, an enterprise software business, surveyed 1,000 Australians from April to May 2020 to capture a snapshot of privacy concerns during the coronavirus pandemic.

In the event of a privacy breach, the potential harm suffered by the affected entity versus the affected individual(s) is disproportionate. The affected entity may incur costs to remedy the breach, suffer penalties (e.g. under the NDB Scheme), and/or take a hit to its reputation and bottom line. These are often short-term economic consequences. The affected individual may suffer longer-term psychological, emotional, financial, or reputational harm.

A 2019 study encouraged clinicians to advocate for a basic right to privacy as a means to safeguard psychological health,[8] after finding that individuals affected by privacy breaches often suffer from anxiety, depression, and post-traumatic stress disorder. The high profile Ashley Madison data breach in 2015 led to two people reportedly committing suicide[9] and many others resigning from their jobs or going through divorces.[10] Identity theft victims reported suffering from stress and anxiety, along with persistent aches and pains.[11]

Comparatively, it's simply not a balancing act.

It's important to consider entities' interests to ensure privacy protections are not disproportionately applied and do not unduly stifle business activities and innovation. But not all entities' interests are created equal and shouldn't be treated as such. For example, journalistic freedoms and healthcare considerations should be balanced with privacy protections, but it's hard to justify balancing privacy protections with an entity's interests relating to its advertising campaigns. Therefore, 'balance' shouldn't be used in this indiscriminate manner; rather, the entity's interests should be considered and weighted according to its nature.

Adding a new object on data minimisation

As CISCO's ex-CEO John Chambers famously remarked, "there are only two types of organisations: those that have been hacked and those that don't know it yet." The first step to reducing the risks of privacy breaches is to eliminate unnecessary collection and storage of data. Therefore, we propose adding a new object on data minimisation: "to promote that entities should collect, use, process, or otherwise handle personal information to the extent adequate, relevant, and limited to what is necessary in relation to the purposes for which they are handled."

[8] Aboujaoude E., Protecting privacy to protect mental health: the new ethical imperative, Journal of Medical Ethics (2019); 45: 604-607.
[9] In 2015, a hacker group stole user information from Ashley Madison, an online website which enabled extramarital affairs. The psychological aftereffects on affected individuals were widely reported. See Australian Broadcasting Corporation, 'Ashley Madison hack may have caused two people to commit suicide, police say' (Webpage) <https://www.abc.net.au/news/2015-08-25/ashley-madison-hack-two-people-may-have-committed-suicide/6721840>. See also The Guardian, 'Life after the Ashley Madison affair' (Webpage) <https://www.theguardian.com/technology/2016/feb/28/what-happened-after-ashley-madison-was-hacked>.
[10] USA Today, 'Anxiety, depression and PTSD: The hidden epidemic of data breaches and cyber crimes' (Webpage) <https://www.usatoday.com/story/tech/conferences/2020/02/21/data-breach-tips-mental-health-toll-depression-anxiety/4763823002/>. Cybercrime and psychiatric experts have commented on the psychological effects on affected individuals resulting from data breaches revealing sensitive, with one even comparing it to the effects of "traditional terrorism".
[11] Identity Theft Resource Centre, 'The Aftermath - The Non-Economic Impacts of Impacts of Identity Theft' (Webpage, 2018) <https://www.idtheftcenter.org/wp-content/uploads/2018/09/ITRC_Aftermath-2018_Web_FINAL.pdf>.

3. Definition of Personal Information

Upgrading the definition of 'personal information'

Current	Proposed
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, or (b) whether the information or opinion is recorded in a material form or not.	Information or an opinion related to an identified individual, or individual who is identifiable: (a) whether the information or opinion is true or not, or (b) whether the information or opinion is recorded in a material form or not. To determine whether an individual is identifiable, account should be taken of all means reasonably likely to be used to identify the individual directly or indirectly.

First, we propose changing the definition of 'personal information' to be information "related to" (rather than "about") an identified individual, or an individual who is identifiable. This expands the reach of privacy protections and removes any confusion of what is personal information.

The use of the term "about" causes unnecessary confusion. From a layperson's view, a lot of information we consider to be personal information may not necessarily be *about* us. Our location data arguably isn't *about* us; it's about our location which reveals things about us. The furniture layout in our homes arguably aren't *about* us,[12] but about our design choices and living density. Video footage from a home security camera isn't *about* a person, but sure reveals things about that person. There are two or more degrees of separation between these forms of information, and information *about* an individual.

At first glance, this may seem unnecessarily argumentative — the information above seems private and should naturally be protected. However, we should ensure the definition is flexible to reduce confusion and mitigate the risks of certain types of information unintentionally not being captured.

Second, we propose tightening the concept of an identifiable individual by removing the reasonableness standard.[13]

- An individual is either identifiable or not; this is binary and shouldn't be subject to reasonableness. The standard for determining whether an individual is identifiable can be subject to reasonableness to ensure it isn't too burdensome (as explored below).
- Having a reasonableness standard adds uncertainty. Entities are currently, in practice, responsible for determining whether an individual is reasonably identifiable (with reference to law and regulatory guidance). This determination is conducted without transparency, and isn't tested or challenged until something goes awry.
- Reasonableness implies a middle-of-the-park judgement, which is too soft and lenient in the light of privacy.

[12] See The New York Times, 'Your Roomba May Be Mapping Your Home, Collecting Data That Could Be Shared' (Webpage, 25 July 2017) <https://www.nytimes.com/2017/07/25/technology/roomba-irobot-data-privacy.html>. iRobot Corporation, the creator of vacuum robot Roomba, was caught in a privacy storm when they reportedly considered using Roombas to map house layouts and sharing that data (including sizes and placements of furniture).

[13] The GDPR doesn't include a reasonableness standard. Under Article 4, paragraph 1 of the GDPR, "personal data is ... any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly…"

Removing the reasonableness standard increases the entities' burden of proving that a person isn't identifiable. So, to ensure that this burden isn't unduly disproportionate, it can be satisfied so long as the entity takes account of all means reasonably likely to be used to identify the individual directly or indirectly. This also encourages more robust anonymisation of information.

Protecting inferred personal information

The issues paper refers to 'inferred personal information' as information collated from a number of sources which reveals something new about an individual. As a general comment, the "reveals something new" language should be reconsidered or clarified. Inferred personal information should cover any inferences, whether it's new or otherwise. For example, if an entity infers my shoe size, it doesn't necessarily reveal something new about me, but it should still fall within this definition.

3. Should the definition of personal information be updated to expressly include inferred personal information?

Yes, it should.

A few years ago, Target started assigning "pregnancy prediction" scores to its female shoppers and used it to promote baby products to expectant mums.[14] A year later, a man walked into Target's Minneapolis store with his teenage daughter to accost the staff for sending coupons for baby products to her. That man later apologised after learning that his daughter was, in fact, pregnant.

This story is commonly cited because it serves as a cautionary tale about how the discovery and use of inferred personal information can erode privacy in a disturbing manner. While using predictive technology[15] isn't the only way to infer personal information, it's possibly the most common and powerful method to do so. Every day, people we've never met are using predictive analytics to discover truths about us (including sensitive ones like our age, sexual preferences, and political views) or predict our future behaviour. These discoveries or predictions are used to inform decisions which affect us in the real world. Most of this happens without our knowledge or consent.

If this gives you pause, you're not alone. In a 2018 Genpact survey,[16] around 71% of respondents didn't want companies using artificial intelligence that threatens to infringe on their privacy (even with their knowledge). More than 50% of respondents were uncomfortable with companies using artificial intelligence to access their personal data. 63% of respondents are worried that artificial intelligence will make decisions that will impact their lives without their knowledge.

[14] The New York Times Magazine, 'How Companies Learn Your Secrets' (Webpage, 16 February 2012) <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.
[15] Predictive technology is a tool which uses machine learning to discover or forecast information, patterns, or behaviours.
[16] Genpact, 'The consumer: Sees AI benefits but still prefers human touch' (Webpage, 2017) <https://www.genpact.com/downloadable-content/the-consumer-sees-ai-benefits-but-still-prefers-the-human-touch.pdf>. Genpact worked with YouGov to conduct a survey of 5,179 people in the United States, United Kingdom, and Australia (1,241 people) in August 2017.

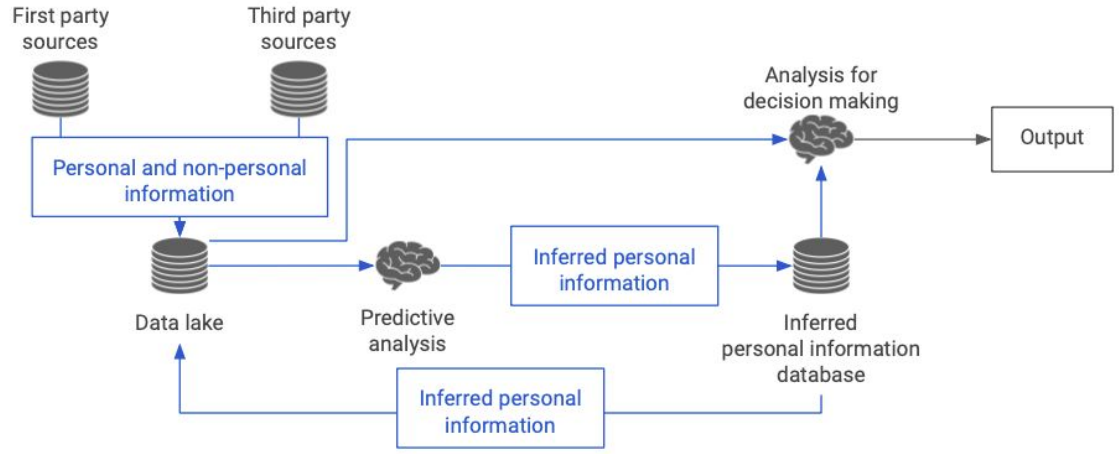

Caption: General flowchart of the discovery, use, and reuse of inferred personal information.[17]

But despite their reticence, discomfort, and worries, entities around the world continue to discover and use inferred personal information in relative obscurity, and are incentivised to do so as it mitigates potential public backlash.

The impacts of decisions made using inferred personal information have real-world consequences that range from mild inconveniences to severe adverse impacts *at scale*. At the milder end of the spectrum, an entity can implement personalised and dynamic pricing by predicting a consumer's demand for a specific product at a certain time. Such practices raise consumer protection concerns, but they may not in themselves justify enhancing privacy protections with respect to inferred personal information.

Caption: Spectrum of potential consequences from the discovery and use of inferred personal information.

However, these protections are needed when we consider that harm can be caused *at scale* by decisions made using inferred personal information. A bank may offer higher loan rates to a certain group of individuals by predicting their repayment abilities based on historical demographic data. This entrenches historical biases, and is discriminatory and manipulative. Predictive policing can reinforce historical racial bias[18] and enable targeted policing of a specific demographic.[19]

[17] For ease, we refer to "non-personal information" as de-identified, anonymised, and pseudonymised information, including information from which an individual is no longer reasonably identifiable, information about households or groups of people, where no one person is reasonably identifiable, or information about deceased individuals. Page 17 of the issues paper describes these forms of information as not being included in the definition of 'personal information'.

[18] Predictive policing uses predictive analytics and other statistical methods to identify potential criminal activity. There is extensive research raising concerns about its effectiveness and how it reinforces biases contained in historical datasets. See, e.g., MIT Technology Review, 'Predictive policing algorithms are racist. They need to be dismantled.' (Webpage, 17 July 2020) <https://www.technologyreview.com/2020/07/17/1005396/predictive-policing-algorithms-racist-dismantled-machine-learning-bias-criminal-justice/>.

[19] See Financial Times Magazine, 'The role of AI in China's crackdown of Uighurs' (Webpage, 11 December 2019) <https://www.ft.com/content/e47b33ce-1add-11ea-97df-cc63de1d73f4>.

Of course, we cannot conveniently dismiss the benefits that predictive technology, and the use of inferred personal information, can offer. Banks could use information on loan repayment abilities to target the delivery of financial education and appropriate financial products (e.g. loans with more practical repayment schedules). Predictive policing data could be used to deploy resources more effectively. Some benefits can be immense (e.g. predictive and personalised medicine).

Ultimately, we need entities to be more disciplined, accountable, and transparent with their discovery and use of inferred personal information. Expressly including inferred personal information as 'personal information' is a good first step towards doing so. This also empowers the OAIC to continue regulating the handling of inferred personal information on an ongoing basis to account for relevant ethical and philosophical implications - e.g. disallowing predictive sentencing.

'Anonymised' information doesn't protect privacy.

> 4. Should there be additional protections in relation to deidentified, anonymised and pseudonymised information? If so, what should these be?

Yes, we support additional protections.

Can information really be anonymised?

In an investigation[20] published as part of the New York Time's *The Privacy Project* series,[21] researchers obtained a set of 'anonymised' mobile phone location tracking data. They used it to isolate a random person and were able to create a personalised diary of their movements. From there, it'll be relatively easy to re-identify that person — noting this is from a single dataset containing only one type of information. Imagine the ease of re-identifying, and creating an accurate profile of, a person using multiple anonymised datasets.

Concerningly, this is not an exception. In one study,[22] researchers published a method which correctly re-identified 99.98% of individuals in an anonymised dataset using only 15 demographic factors. Journalists in Germany re-identified politicians in an anonymised set of 3 million German citizens, which revealed their sensitive information like medical information and sexual preferences.[23] Closer to home, researchers re-identified anonymised health data released by the Australian Department of Health in only 6 weeks by linking unencrypted parts of the record with known information about the individuals.[24] Precise location data is arguably impossible to anonymise.[25]

Given the ease of re-identification, anonymising information doesn't protect privacy anymore. The Privacy Act currently fails to recognise this. To enhance privacy protections, the Privacy Act has to implement appropriate parameters for activities involving non-personal information.

---

[20] The New York Times, 'Twelve Million Phones, One Dataset, Zero Privacy' (Webpage, 19 December 2019) <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>. The New York Times obtained a file containing logs of the movements of people with mobile phones, holding more than 50 billion location pings from phones of over 12 million Americans. By analysing this dataset, the New York Times' researchers mapped movements in the Pentagon, White House, to and from celebrities' homes, among other things.

[21] The New York Times, 'The Privacy Project' (Webpage) <https://www.nytimes.com/series/new-york-times-privacy-project>. This project's an initiative by the New York Times to explore technology, its impacts and how we can control it. It's organised into three phases: education, debate, and calls to action.

[22] Rocher, L., Hendrickx, J.M. & de Montjoye, Y. Estimating the success of re-identifications in incomplete datasets using generative models. *Nat Commun* 10, 3069 (2019); https://doi.org/10.1038/s41467-019-10933-3 (**Rocher Paper**).

[23] Hern, A. 'Anonymous' browsing data can be easily exposed, researchers reveal. The Guardian (1 Aug 2017), cited in the Rocher Paper.

[24] Culnane, C., Rubinstein, B. I. P. & Teague, V. Health data in an open world. https://arxiv.org/abs/1712.05627, cited in the Rocher Paper.

[25] See footnote 20.

Protections for non-personal information[26] sources

Non-personal information can be re-identified in two main ways: (a) by combining it with personal information, or (b) by combining it with other non-personal information.

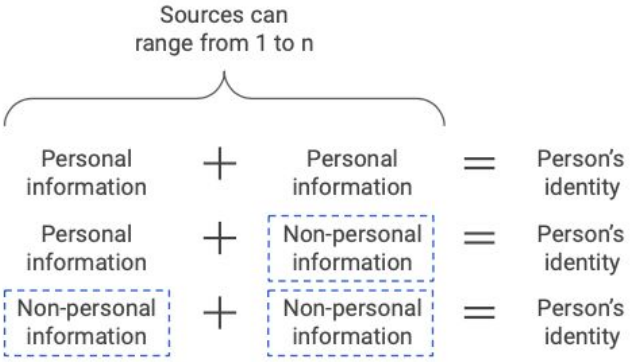

Caption: General illustration of how re-identification may occur.

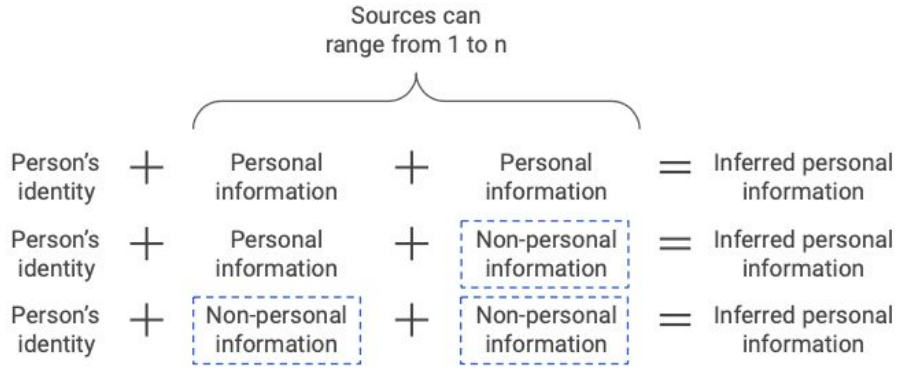

Caption: Re-identification can occur indirectly as part of the process to infer personal information.

We propose three general additional protections for non-personal information: (a) securing Valuable NPI, (b) regulating the use of non-personal information, and (c) regulating data brokerage. "**Valuable NPI**" refers to one or more sets of non-personal information (as a whole) which are used or can be used (taking into account all reasonably likely means) to infer personal information or re-identify the individual to which that information relates.

Securing Valuable NPI

First, Valuable NPI actually used in re-identification should be segregated from other personal information sources or Valuable NPI that were also used. Further, the applicable method (e.g. algorithm, analytical process) of re-identification, or that resulted in re-identification, should also be segregated. Separating the components mitigates the risks of unauthorised or malicious actors 'recreating' the inference or re-identification process.

Second, depending on its nature, Valuable NPI should also be subject to similar security requirements as personal information — e.g. protected from misuse, interference and loss, as well as unauthorised access, modification or disclosure. Anonymised mobile phone location data will fall within this category. This wouldn't be excessively onerous as most entities should have established security practices, and information can be encrypted at rest and in transit at low costs.

[26] See footnote 18.

However, to ensure proportionality, these security requirements shouldn't apply to information that is publicly available. By using something as innocuous as Facebook likes only,[27] researchers could predict highly sensitive things about a person — including their age, race, and political views — with surprising accuracy. If an entity de-identifies those Facebook likes (i.e. segregates the personal information from the Valuable NPI), then those Facebook likes are exempt from the security requirements provided they are publicly available (e.g. likes of public pages).

Regulating the use of non-personal information

We propose the development of lawful bases for which non-personal information can be used to infer personal information or for re-identification purposes. For obvious reasons, obtaining consent to use non-personal information is not possible.

From a commercial standpoint, there are many applications of using non-personal information to infer personal information. For example, a major supermarket chain runs a rewards program that collects activity data like grocery spend. It anonymises this information and sells it to a ridesharing app. The ridesharing app combines the anonymised information with other information it's collected to build a more complete profile of its users. It then uses these insights to provide differential, personalised pricing to maximise each user's engagement and value.

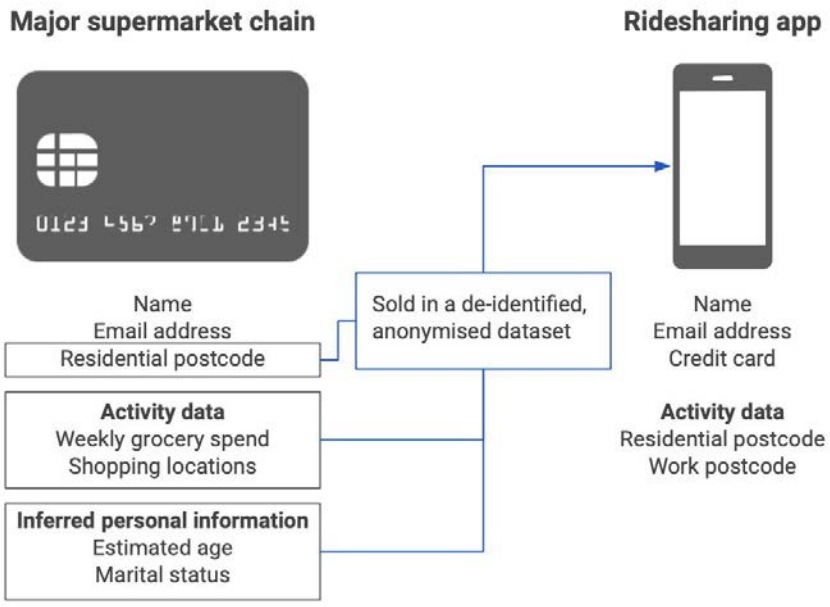

Caption: Example of anonymised data being shared between two businesses.

As these activities have real-world impacts, there should be a manner by which entities keep the affected individuals informed that re-identification has occurred. Those individuals should also be able to exercise a level of control over such activities, like having the ability to opt-out.

From a public interest standpoint, entities should be able to use non-personal information to infer personal information or for re-identification purposes for similar lawful reasons as personal information under APP 6 — e.g. enforcement related activities, relating to health situations, or authorised by Australian law.

[27] Kosinski, M., Stillwell, D., and Graepel, T., Private traits and attributes are predictable from digital records of human behavior (2013) PNAS April 9, 2013 110 (15) 5802-5805; https://www.pnas.org/content/110/15/5802. In a study conducted with over 58,000 volunteers, researchers found that Facebook likes can automatically and accurately predict a range of highly sensitive personal attributes including: sexual orientation, ethnicity, religious and political views, personality traits, intelligence, happiness, use of addictive substances, parental separation, age, and gender.

Regulating data brokerage

Data brokerage is the business of buying and selling data. The data brokerage industry is reported to be worth US$200 billion. Because this industry operates in relative obscurity, it's hard to determine from a desktop analysis the exact extent and nature of data that these businesses trade. Data brokers have, for example, reportedly sold location data which helped predict the race, age and gender of *Black Lives Matter* protestors, and helped immigration authorities catch undocumented immigrants.[28] So, it's safe to assume that these businesses have undue control of and access to information (both personal and non-personal).

As part of this Review, we have an opportunity to implement standards that require greater transparency into data brokerage practices, and measures that enable lawmakers and the public to assess the effectiveness of privacy protections on those practices. One way of doing so could be to establish a data-broker clearinghouse[29].

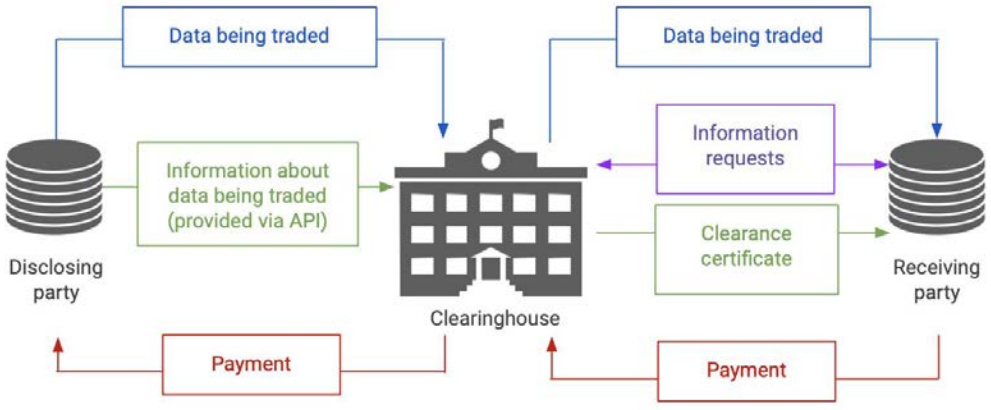

Caption: Illustration of data brokerage clearinghouse.

The clearinghouse will be a designated intermediary supervised by the OAIC, ACCC, or such other appropriate regulatory body. Any data being traded for commercial purposes will flow through it. Generally, the clearinghouse will:

- establish a publicly available set of rules governing data trades — e.g. provenance of data, legitimacy of proposed uses, and types of data that cannot be traded;
- collect relevant information (in a standardised manner) from the disclosing party via an API; and
- be able to request additional information from the receiving party if the disclosing party's diligence is insufficient.

These clearinghouses will need to be appropriately licensed by the OAIC, and subject to regular reporting and audit requirements, not unlike the regulatory structure for clearinghouses within the Australian financial system. The OAIC (and interested persons) should be wary of over-engineering the rules governing data-broker clearinghouses, at least initially, given they wouldn't yet be fundamental to the economic stability of Australia. In parallel, lawmakers can also require data

[28] See TechCrunch, 'Data brokers track everywhere you go, but their days may be numbered' (Webpage, 9 July 2020) <https://techcrunch.com/2020/07/09/ data-brokers-tracking/>, which noted Google and Apple's efforts to reduce the power of data brokers by allowing users to opt-out of ad tracking and, for Apple devices, opt-out of app tracking, too.
[29] See Time, 'You Deserve Privacy Online. Here's How You Could Actually Get It' (Webpage, 2019) <https://time.com/collection/davos-2019/5502591/tim-cook-data-privacy/>. In this op-ed by Tim Cook, he suggested that the United States Federal Trade Commission establish a data-broker clearinghouse.

brokers to register with a central registry and provide sufficient information to enable the relevant regulator to maintain effective oversight of the industry.[30]

Of course, these proposals will increase regulatory burden and the regulator will need to be appropriately resourced. However, we believe these additional efforts and resources are justified as the data brokerage industry is in sore need of greater regulation and supervision.

Greater clarity for 'personal information'

> 5. Are any other changes required to the Act to provide greater clarity around what information is 'personal information'?

Yes. The definition of 'personal information' should clearly include (as a non-exhaustive list):

- online identifiers (as described in the issues paper), internet protocol (IP) addresses, crash reports, system activity, referrer URL of requests, data about interaction with apps, browser and device types, application, carrier name, metadata, and operating system; and
- activity information like networks and connections, purchase activity, time, frequency and duration of an activity, browsing history, routing information, and privacy settings.

4. Small business exemption

We propose removing the small business exemption in support of the reasons cited in the issues paper, e.g. the recommendation in ALRC Report 108. Additionally:

- **The small business turnover test is inappropriate.** Turnover isn't an accurate proxy for a business' potential impact on privacy-related issues. A data analytics startup backed by private capital can develop a product using extensive amounts of personal information without having to generate revenue.

  In 2018, SmartCompany[31] analysed the privacy posture of the top 100 fastest growing businesses in Australia, which included 54 businesses with annual turnovers of under $10 million. Of the 100 businesses, 44% of them didn't appear to comply with Australian privacy laws. While it's unclear what proportion of these non-compliant businesses are small businesses, the offending businesses are in the most data-driven sectors. The notion that, in 2018, businesses in finance and insurance, education and training, and advertising industries are still handling personal information in non-compliant manners should call for concern.
- **Costs to comply aren't high anymore.** Collecting and storing personal information offline is getting rarer. Small businesses today use a range of software products and cloud services. These software services (e.g. Xero, Stripe) often have in-built security functions. Cloud

[30] The state of Vermont in the United States enacted the country's first legislation requiring data brokers to register with its state secretary. Vermont's Act No. 171 (2018), available at <https://legislature.vermont.gov/Documents/2018/Docs/ACTS/ACT171/ACT171%20As%20Enacted.pdf>, defines data brokers as businesses which aggregate and sell data about consumers with whom the business doesn't have a direct relationship. Vermont's 9 V.S.A. § 2446, available at <https://legislature.vermont.gov/statutes/section/09/062/02446>, subjects data brokers to a registration requirement. Further, Vermont's 9 V.S.A. § 2447, available at <https://legislature.vermont.gov/statutes/section/09/062/02447>, subjects data brokers to privacy protection requirements.

[31] SmartCompany, 'A worrying number of Australian companies don't comply with privacy laws or have secure websites' (Webpage, 6 July 2018) <https://www.smartcompany.com.au/technology/australian-companies-privacy-laws-secure-websites/>. The analysis was conducted on the Australian Financial Review's 100 fastest growing companies in 2017, comprising 54 small companies (annual revenue under $10 million), 45 medium companies (annual revenue between $10 million and $249 million), and 1 large company (annual revenue over $249 million). The smallest company had an annual revenue of just over $1.5 million.

services like Amazon Web Services and Microsoft Azure offer encryption tools at low, variable rates. Additionally, they provide accessible how-to guides on building security into a business' cloud environment.

- **Businesses should be built with privacy at its core.** All businesses should practise good privacy hygiene no matter its size. In any case, a proportion of small businesses already opt-in to comply with the APPs.

5. Notice of Collection of Personal Information

> 20. Does notice help people to understand and manage their personal information?
> 25. Would a standardised framework of notice, such as standard words or icons, be effective in assisting consumers to understand how entities are using their personal information?

Reduce barriers to reading privacy notices

Put simply, we have to reduce barriers to people reading privacy notices[32]. The 2020 ACAP survey identified length and complexity as the main deterrent. So, addressing those issues is a sensible starting point to increase readership.

We should also bear in mind that the 2020 ACAP survey was conducted on individuals aged 18 and above. That is, the survey data doesn't cover individuals aged below 18, who are most unlikely to read or understand a privacy notice. Granted, such individuals do not have legal capacity to give consent. However, in practice, it's inevitable that they are providing personal information without an adult looking over their shoulders. A large (and increasing) number of them are online. 79% of individuals in Year 12 or below (over 7.06 million) had access to the internet in 2016 — 2017,[33] largely for banking, social networking, and entertainment. So, we have to design a better way to encourage them to read privacy notices, and make it easily understandable.

We propose requiring each entity to provide, in an easily accessible manner, a visual privacy summary to complement its privacy notice. This is consistent with the ACCC's suggestion in the DPI Report,[34] as noted in the issues paper.

Visual privacy summaries

The visual privacy summary initiative can be implemented as follows:

- The OAIC will design a set of standardised privacy icons in collaboration with interested persons. These icons will be managed by the OAIC.
- Each privacy icon will have two components: a fixed element that represents the relevant category, and a variable element that reflects the entity's position on that category. For example, there could be a three-way crossroads icon for offshoring. Entities will then have to select the one which reflects their position e.g. Australia only, offshoring to countries with comparable privacy regimes, and offshoring to countries relying on contractual privacy regimes.

[32] The term "privacy notice" is not used consistently and may sometimes refer to privacy policy summaries or updates. For ease, we refer to "privacy notices" in this submission in an interchangeable manner with "privacy policies".
[33] ABS, 'Household use of information technology (2016-2017)' (Webpage) <https://www.abs.gov.au/statistics/industry/technology-and-innovation/household-use-information-technology/latest-release>.

[34] ACCC, Digital Platforms Inquiry (n 1) 403.

For illustrative purposes:

No offshoring. Australia only.

Offshoring to countries with comparable privacy laws

Offshoring to countries relying on contracts to ensure privacy

- The privacy icons must be displayed in a predetermined order. This ensures consistency and optimises the ability for individuals to recognise.
- The visual privacy summary must be made available and displayed in an easily accessible way. It must also link to the full-form privacy notice.

For reference, in 2011, Mozilla proposed a set of standard privacy icons which can be used in privacy notices.[35] This project never achieved mass adoption, but is a great representation of how graphics can be used to make complex information more accessible.

There are four main reasons for this proposal:

- **Simplicity should not come at the cost of detail.** The information in privacy notices can be complex. But sometimes it's for good reason. We'd probably want in-depth explanations about how the advanced analytics or tracking tools used on a website with some technical accuracy. The visual privacy summary adds a layer of simplicity without getting rid of the detail (i.e. the privacy notice isn't replaced).
- **Privacy notices can be ambiguous.** They use the term "may" a lot, which is appropriate given activities described in privacy notices are often discretionary. However, ambiguity is antithetical to clarity. Further, as ambiguity in privacy notices have become generally accepted, people are dissuaded from reading it. After all, every entity may collect, use, and disclose our personal information.

  The visual privacy summary provides an additional level of certainty by using standardised privacy icons and a predetermined layout. Readers can easily identify that there are fixed categories of privacy protections, and fixed ways to think about them.
- **Standardisation reduces complexity.** Privacy notices tend to be bespoke — from the way it's drafted, to the format and style. Having new ways to build readable privacy policies helps.[36] But with many options being used, and a slow pace of adoption, we are still faced with comparability issues. It's still challenging to read and compare privacy notices to what we have seen before, even for an experienced person.

  Using standardised privacy icons and a predetermined layout enables individuals to quickly familiarise with the key concepts. There is also greater incentive to understand these icons as it that knowledge will be translatable across all websites and apps.
- **Graphics can be universally understood and more digestible.** We process images faster than reading and processing words. In our age of distraction and short attention spans, visually representing privacy notices (or key concepts in them) is the optimal way to enhance

[35] See Mozilla, 'Privacy Icons Project (beta)' (Webpage) <https://wiki.mozilla.org/Privacy_Icons>.
[36] See, e.g., Juro, 'Privacy by design: building a privacy policy people actually want to read' (Webpage, 5 May 2018) <https://medium.com/juro-blog/privacy-by-design-building-a-privacy-policy-people-actually-want-to-read-bd8eb4a52403>.

engagement and understanding. Individuals under 18 will also likely understand a graphical representation of privacy concepts better.

Of course, the privacy icons have to be carefully designed. They need to be intuitive and fairly self-explanatory. We recommend that the OAIC maintains sole control of the icons and collaborates with interested persons in designing it.

The proposal and comments in this section also apply to questions 24 and 30.

6. Consent to collection and use and disclosure of personal information

27. What approaches should be considered to ensure that consent to the collection, use and disclosure of information is freely given and informed?

We support the ACCC's Recommendation 16(c) relating to consent.

32. Should entities collecting, using and disclosing personal information be required to implement pro-privacy defaults for certain uses and disclosures of personal information?

We support the ACCC's Recommendation 16(c) to strengthen pro-consumer defaults. Specifically, default settings with respect to handling of personal information for a purpose other than the performance of a contract should be preselected to 'off', as preferred by a majority of digital platform users.

Generally, entities operating digital services have disproportionate power over individuals who use, or wish to use, those services. The entity may argue that individuals dissatisfied with its privacy posture could always choose not to use their services. However, this argument is impractical - we can't just 'choose' not to use Facebook, Google Search, YouTube, or WhatsApp - and ironic, given these entities devote their time to designing and providing products that become essential in our world.

Default pro-individual protections can be found in other areas of the law, particularly for consumers who may otherwise be unable to protect their own interests. For example, in New South Wales, buyers of off-the-plan apartments from developers are entitled to a cooling-off period of 10 business days by default, which can only be expressly waived after receiving legal advice. This cooling-off period enables the buyer to take a breather to assess the purchase and their affairs, outside of the pressures to sign and lock-in a "one-off" great deal.

We think that pro-privacy defaults would help re-balance the power equation, and protect those who are unwilling or unable to assess the extent of personal information handling that they're being asked to consent to (e.g. due to consent fatigue or plain convenience). Additionally, we propose the following protections:

- Only the individual can opt-out or change the settings, and such a change will only be valid if the entity can demonstrate that it has clearly explained how the changes affect that individual - e.g. how personal information will be handled upon changing that setting and how this would compare to the usage of their information if the pro-privacy settings were retained.
- The individual's access and use of the services should not be unfairly or unreasonably reduced or hampered because they rely on the pro-privacy defaults. That is, the entity must

not engage in conduct that unduly influences the individual to opt-out of those pro-privacy defaults.

Consent from children

33. Should specific requirements be introduced in relation to how entities seek consent from children?

Yes. But it's not a simple challenge to tackle.

The OAIC describes four characteristics to determine whether a person has capacity to give consent, being the person: understands that they're being asked to give or not give consent, understands the consequences of giving or not giving consent, based their decision on reason, and can communicate their decision.[37] There's a blanket exception for minors which, in Australia, is generally people under 18.[38] This exception makes two inaccurate presumptions:

- Minors don't, or can't, meet the characteristics above. Arguably, a 16 or 17 year old could have the same comprehension of consent and privacy consequences, particularly given the amount of time teenagers spend on the internet today.
- Minors will ensure appropriate consent is given. But practically, can you imagine children or teenagers reading through privacy notices with their parents or guardian before using a website or app? Probably not.

We need to align consent requirements for children with the practical realities. To do so, we'll need to consider two main things.

Can a child give consent? The ALRC Report 108 considered a few policy approaches for assessing the capacity of minors,[39] and noted that children between 14 and 16 may have the cognitive ability to make independent decisions but still remain more susceptible to psychosocial factors (which are dependent on the circumstances). The report also suggested that it may be appropriate to make the minimum age for consent dependent on the nature of the personal information — i.e. consent relating to sensitive information has to be given by an adult.

For children under 16, we could consider a test similar to the Gillick competency.[40] This test, which originated from English medical law and has been adopted in Australia,[41] is used by the courts and health professionals to identify children under 16 who have the legal competence to consent to certain medical treatments. It takes into account four things:

- maturity;
- the child's experiences and ability to manage influences on their decision making (e.g. information, peer pressure, fear);

[37] OAIC, 'Consent to the handling of personal information' (Webpage) <https://www.oaic.gov.au/privacy/your-privacy-rights/your-personal-information/consent-to-the-handling-of-personal-information/#CapacityToConsent>. The OAIC provides guidance on express consent, implied consent, bundled consent, and what consent involves - i.e. informed, voluntary, current and specific, and capacity.
[38] The United Nations Convention on Children's Rights defines a child as any person under 18.
[39] ALRC, 'Possible models for assessing capacity', *For Your Information: Australian Privacy Law and Practice (ALRC Report 108)* (Webpage, 17 August 2010) <https://www.alrc.gov.au/publication/for-your-information-australian-privacy-law-and-practice-alrc-report-108/68-decision-making-by-and-for-individuals-under-the-age-of-18/possible-models-for-assessing-capacity/>.
[40] *Gillick v West Norfolk and Wisbech* [1986] AC 112.
[41] *Secretary, Department of Health and Community Services v JWB and SMB* (1992) 175 CLR 218.

- intelligence; and
- the child's understanding, and ability to weigh risk and benefit, and consideration of longer-term factors like effects on family life and schooling.

Privacy is a complex subject and there are potentially long-term consequences to giving relevant consent (even when taking into account a right to erasure, if one is recognised). Therefore, the child will need to have a high level of competence. It may be that, after extensive consideration, we find that children don't have competence to give consent. But it is a necessary exercise to undertake.

What is the best way to inform children? We'll defer to pedagogic and child development experts on this subject. In our view, there are two challenges that require further thought: what is a sufficient level of information that a child needs to understand in order to give informed consent, and how will an entity demonstrate that it's satisfied that a child understands. In terms of delivering the information, implementing visual privacy summaries (as explored above) could encourage engagement and improve understanding.

Privacy protections relating to IoT devices

> 34. How can the personal information of individuals be protected where IoT devices collect personal information from multiple individuals?
> 43. Are the security requirements under the Act reasonable and appropriate to protect the personal information of individuals?

The Australian Strategic Policy Institute (**ASPI**) remarked in its *Internet of Insecure Things* report[42] that security certifications and security regulations should be proportionate to the risk profile of the IoT device. This is consistent with APP 11 — if an IoT device scales up in risk (e.g. starts collecting more sensitive information), it'll be required to step up its security measures.

However, the ASPI's message in that report supports the notion that there should be no circumstance whereby security measures can be excused. ASPI highlights the threat of security vulnerabilities in IoT on critical infrastructure. There are a growing number of distributed denial of service (**DDoS**) attacks which exploit IoT vulnerabilities. Affected IoT devices may be 'low risk' like a smart coffee machine, but when exploited with millions of other low risk IoT devices, the combined impact is enormous. This may not directly impact privacy *per se*, but definitely affects the resiliency of security measures built to protect privacy.

So, we need better security discipline around IoT and any other technologies that have network-reach impacts. Tightening APP 11 may be the right starting point.

APP 11 requires entities to take reasonable steps to protect personal information from misuse, interference, and loss, and unauthorised access, modification or disclosure. Whether an IoT device collects personal information from one or multiple individuals, the entity which handles the data should have security measures built in that meet the standards in APP 11. Rather, a better question is whether the standards in APP 11 are sufficient for IoT — in particular, "reasonable steps".

[42] Australia Strategic Policy Institute, 'The Internet of Insecure Things' (Webpage, 2018) <https://www.aspi.org.au/report/InternetOfInsecureThings>.

Tightening up “reasonable steps”

The OAIC provides guidance that reasonableness depends on the circumstances,[43] like the entity’s nature (e.g. resources and size), the amount and sensitivity of personal information held, possible adverse consequences in a breach, practical implications of implementing the security measures (e.g. time and cost), and whether the security measure is privacy invasive.

We largely agree with the elements of reasonableness, as they require proportionate ‘step ups’ of security measures based on risk profile. However, we’re concerned with the ‘practical implications’ element. The OAIC explains that taking into account practical implications doesn’t mean that entities are excused from taking particular steps only if it is inconvenient, time-consuming or imposes some cost. But entities may be excused from doing so if the burden is excessive in all the circumstances.

Given the importance of privacy, there should simply be no excuses at all. An entity’s size or resources should not underlie what ‘reasonable steps’ it should take, particularly when (as explored above) a small business could in theory have limited resources but conduct data processing with a profound privacy impact. If, in the circumstances, an entity determines that implementing the security measure is excessively burdensome (e.g. hugely inconvenient, disproportionately time-consuming, and very expensive), then the entity should not collect that personal information. It may require the entity to rethink its business model or product and, in some cases, slow innovation. But slow isn’t always bad especially when it comes to privacy; taking a measured approach is the right thing to do.

Right to erasure

> 46. Should a ‘right to erasure’ be introduced into the Act? If so, what should be the key features of such a right? What would be the financial impact on entities?

We support a ‘right to erasure’ that’s based on Article 17 of the GDPR. Regarding the financial impact, we’ll defer to organisations that have implemented relevant systems and processes to comply with this requirement under the GDPR.

Direct right of action

> 56. How should any direct right of action under the Act be framed so as to give individuals greater control over their personal information and provide additional incentive for APP entities to comply with their obligations while balancing the need to appropriately direct court resources?

We strongly support a direct right of action under the Privacy Act for the reasons outlined by the ACCC in the DPI Report.

The argument that introducing a direct right of action would ‘open the floodgates’ to court proceedings is not a persuasive basis to prevent an individual from seeking judicial resolution of their dispute. As raised by the ACCC in the DPI Report, judicial interpretation assists better understanding and application of the APPs, particularly given their principles-based nature. Also, there are sufficient safeguards in the existing judicial process to prevent trivial, frivolous, or irrational litigation from having its day in court, including:

[43] OAIC, ‘Chapter 11 - APP 11 - Security of Personal Information’ (Webpage, 22 July 2019) <https://www.oaic.gov.au/privacy/australian-privacy-principles-guidelines/chapter-11-app-11-security-of-personal-information/>.

- a lawyer not being able to act unless they believe their client has reasonable prospects of success;
- disincentives for instituting (or continuing) litigation like the significant associated costs, the risks of reaching the end only to have an adverse costs order (e.g. having to pay their own legal costs and potentially the costs of their opponent), and the impact of Calderbank offers on costs orders which incentivise out-of-court settlements; and
- particularly with respect to self-represented litigants, the risk of being declared as a vexatious litigant, thereby making it harder for them to institute proceedings in the future.

To ensure that court resources are appropriately and proportionately allocated, we think that the direct right of action should be framed as follows:

- any individual can bring a direct right of action for any interference with privacy, whether trivial or not. However, vexatious or frivolous claims shouldn't be entertained. Of course, the damages recoverable by an aggrieved person will be proportionate to the nature of the interference (e.g. severity and harm suffered);
- the individual alleging an interference with privacy can elect to commence legal proceedings in court or undergo an alternative process (e.g. conciliation); and
- the OAIC can, or can elect a qualified representative to, conciliate disputes relating to an interference with privacy. Conciliation provides disputing parties with a flexible, confidential, speedier, and potentially more cost-effective way to resolve the dispute. It'll also reduce the burden on the courts.

7. Conclusion

We appreciate, and thank the Australian Government for, the opportunity to comment on the issues under consideration in the Review and engage in this conversation. This submission is written by the named individuals only and does not purport to express the views or opinions of any other person.

- ETX -