Submission to Inform Government Response to Privacy Act Review Report

School of Computing and Information Systems

The University of Melbourne

The School of Computing and Information Systems (CIS) welcomes the Report into the Review of the Privacy Act and opportunity to comment on the proposals contained within.

The Privacy Act review presents a critical opportunity to address long-standing deficiencies in Australia's privacy laws, and undesirable business practices more broadly. We believe important amendments are required for the Privacy Act Report's proposals to address privacy concerns and protection of personal information.

Executive Summary of Recommendations

- **Proposal 4.4:** We argue that the definition of 'personal information' should extend to the concept of individuation. Doing so would be consistent with community values of respecting individual's autonomy and will not unduly limit valuable uses of data where individuals have provided substantive consent to such uses.
    - Identification and individualisation should not be the only measure when performing risk assessment of de-identified data or determining the level of de-identification. Harm to an individual can be caused even if they are not "singled out", that is, they can be targeted if narrowed down to a few individuals.[1]
- **Proposals 4.4:** We note that the report's understanding of individuation is fundamentally flawed, and support the definition articulated by Anna Johnston of Salinger Privacy: '*the ability to disambiguate or 'single out' a person in the crowd, such that they could, at an individual level, be tracked, profiled, targeted, contacted, or subject to a decision or action that impacts upon them - even if that individual's 'identity' is not known (or knowable)*'
    - We further submit that there is a meaningful distinction to be maintained between uses of the words 'distinguishable' and 'identifiable' in this context. 'Distinguishable' should be understood in the sense of being able to 'single out', consistent with Anna Johnston's definition of individuation, whereas 'identifiable' should be understood as being able to distinguish someone by virtue of their natural or legal identity.
- **Proposal 4.4:** We do not support the use of functional de-identification, as it is an ineffective way of protecting individuals' privacy in data sets. This is because individuals can be re-identified through other means, such as combining the functionally de-identified data with other available data. Allowing functional de-identification will permit APP entities to retain insecure data, and it

[1] Alexandra Wood et al, 'Differential Privacy: A Primer for a Non-Technical Audience' (2018) 21(17) *Vanderbilt Journal of Entertainment & Technology Law* 209.

is not possible to distinguish identifying and non-identifying information as implied by the review.

- **Proposal 4.5:** The proposal to amend the definition of de-identification still does not address the fundamental issues with de-identification in the Privacy Act; de-identification is a critical issue that warrants special attention. We note several issues with a definition of de-identification:
    - The discussion in section 4.5 of the report does not address the purpose of de-identification, which may provide a loophole by which APP entities can employ de-identification to avoid procuring consent. This would unacceptably undermine individual autonomy.
    - De-identification is not a risk-free exercise, and nor does it guarantee that re-identification will not happen. That said, there exists effective technologies for de-identification, which are effective for their clear definition of privacy properties to be preserved in respect of an identified threat model.
    - Any references to 'reasonableness' in relation to *re*-identification must also account for best practice. We are concerned that reference to '*… [not] reasonably identifiable in the **current context'*** may preclude consideration of the fact that data may exist *in the future* (but not in the 'current context') to make de-identified data re-identifiable. To ensure this is avoided, statutory references to 'reasonable' should also mandate consideration of 'best practice'.
- **Proposals 4.6, 4.7:** We strongly oppose criminalising re-identification, for it is ineffective at protecting privacy because a) re-identification may not be detectable, b) the law may not apply or be enforceable outside Australian jurisdictions, c) it would stifle whistle-blowers.
    - We suggest that inadequate de-identification is the real problem this proposal seeks to address, and thus refer to our comments in relation to proposals 4.4–4.5.
- **Proposal 4.8**: We have several concerns in relation to this proposal:
    - It is not clear whether allowing re-identification of de-identified data supplied from an individual undermines APP 2.
    - Excluding re-identification of de-identified data by an APP entity appears to undermine APP 11.3.
    - No explanation is offered as to how such data could be treated like other data.
- **Proposals 10.1 – 10.3:** We support the introduction of mandatory, pro-privacy default settings to ensure consumers are adequately protected.
    - While we support improving the substance and presentation of privacy policies, however this does not justify reliance on formative consent of individuals, who cannot be expected to read long and dense privacy policies, as numerous academic works have affirmed.
- **Proposals 14.1 – 14.3:** Considering our comments on Proposal 4.8, we submit that any exemption for researchers should not be overly broad, as this may become permissive of re-identification for the purpose of research on data subjects.
    - The presently stated exemption may also prohibit security consultants from performing audits of de-identification practices.
- **Proposal 20.2:** Provide individuals with an unqualified right to opt-out of their personal information being used or disclosed for direct marketing purposes. Like existing requirements under the Act, entities would still be able to collect personal information for direct marketing without consent, provided it is not sensitive information, and the individual can opt out.

    - It is not clear why this proposal states sensitive data and not personal as it should include both. Users should be given a form to consent on *any* information about them to be used for marketing or not.
- **Proposals 21.2, 21.3, 21.5:** We are broadly supportive of strengthening the OAIC guidelines, in particular enhancing guidance to more clearly articulate non-exhaustive baseline security requirements and steps to ensure protection and secure deletion of data.
    - We emphasise though that these articulated minima *must not* supplant the requirement of ensuring measures are reasonable in the context in which they are taken.
- **Proposal 21.1:** Proposing 'reasonable steps' incorporate technical and organisation measures is unlikely to have much effect, as organisations appreciate that data protection is a matter of both governance and technical controls.
- **Proposal 26.1:** We support the introduction of a direct right of action (DRA), but do not support it being subject to a complaint gateway. This obviates the stated purpose of the DRA by imposing an unnecessary administrative burden on claimants and does not strike the right balance between preserving court resources and allowing individuals to assert their rights in a court of law.

Personal information, de-identification and sensitive information

Proposal 4.4 'Reasonably identifiable' should be supported by a non-exhaustive list of circumstances to which APP entities will be expected to have regard in their assessment.

We disagree with the report's conclusion that personal information should *not* extend to the concept individuation, as discussed in section 4.3.4 of the report under proposal 4.4.

This conclusion is partly founded on an erroneous understanding of the concept. Furthermore, not extending personal information to the concept of individuation is inconsistent with community values of respect for individual autonomy. Valuable uses of data will *not* be unjustifiably limited, where individuals have given proper substantive consent to those uses.[2]

The correct understanding of individuation

The report correctly defines individuation when quoting Anna Johnston (Salinger Privacy) as "the ability to disambiguate or 'single out' a person in the crowd, such that they could, at an individual level, be tracked, profiled, targeted, contacted, or subject to a decision or action that impacts upon them - even if that individual's 'identity' is not known (or knowable)"[3]. The report then proceeds to offer a paraphrase that demonstrates an erroneous understanding of the definition:

> *"The concept of individuation as understood by the Review is where information relating to an individual reveals their characteristics and can be used to impact them even though they are not reasonably distinguishable or distinguishable from all others. In the context of targeted content and advertising, information relating to an individual 'individuates' them*

[2] See our comments to proposals 10.1 – 10.3 of our submission.

[3] Anna Johnston, 'Individuation: re-imagining data privacy laws to protect against digital harms' (Working Paper Vol 6 No 24, Brussels Privacy Hub, July 2020).

> *from others and can be used to target them even though they are not reasonably identifiable."*[4]

The first sentence is incorrect: individuation is using information about an individual to distinguish them even though they may not be identifiable in the traditional sense, due to the absence of data pertaining to their legal or natural identity, such as their name. It is expressly about distinguishing individuals even though you may not know their common identity, since distinction is all that is required to target at an individual level. The second sentence however is inconsistent with the first by suggesting the individual is individuated from others and allows them to be targeted, which implies distinction.

The report's definition also conflates 'distinguishable' with 'identifiable'. Distinguishable should be understood as being able to tell people apart, whereas identifiable should be understood as being able to identify an individual by their legal or natural identity.

'Personal information' should incorporate individuation.

The review concludes *"Extending the definition of personal information in this way could unjustifiably limit valuable uses of data in ways which do not harm or affect the individuated person"*. This position is not adequately justified.

Firstly, the review's stance diminishes individual autonomy. In line with the object and purpose of the act, we argue that individual should decide what is in their own interest, and not the APP entity. It should be noted that the review was triggered because of the ACCC Digital Platforms Inquiry, whose report stated:

> *"The ACCC considers that the Privacy Act needs reform in order to ensure consumers are adequately informed, empowered and protected, as to how their data is being used and collected."*[5]

Furthermore, the inclusion of individuation would not prohibit valuable uses of data, subject to the individual providing substantive consent. On the matter of substantive consent, we refer you to our submissions in respect of proposals 10.1 – 10.3.

Functional de-identification weakens privacy protections for individuals

The review states:

> *"APP entities may be able to engage in 'functional de-identification' with strict organisational and technical controls so that identifying information is separated. This enables risk managed use, even though without the controls the information would be personal information."*

We are concerned that allowed functional de-identification would weaken the Act in comparison to what the Act currently requires. The proposed changes would not require an APP entity to ensure indistinguishability. It would merely require a form of data masking that was functionality controlled. As such, APP entities would be able to retain even more data about individuals having not had to apply any basic indistinguishable methods.

---

[4] Page 26 of the report.

[5] ACCC, *Digital Platforms Inquiry Final Report* (Report, June 2019) 3

This promotes a common misconception about de-identification through separating attributes. While attributes that clearly identify individuals may be stripped, remaining attributes may be combined with one another or may be linked to external data, to identify individuals.[6] It is not possible to distinguish identifying and non-identifying information as implied by the review.[7]

Proposal 4.5 *Amend the definition of 'de-identified' to make it clear that de-identification is a process, informed by best available practice, applied to personal information which involves treating it in such a way such that no individual is identified or reasonably identifiable in the current context.*

This proposal does not address the fundamental issues with *de-identification* within the Privacy Act—a critical issue with Act. In its current form, the Act requires protections for data *unless* it is de-identified but provides no falsifiable definition of when data is adequately de-identified.

As a result, it is possible that de-identified data may be released (and therefore exempted from the privacy protections required of APP entities for personal or sensitive information) yet still contain identifiable individuals. The issue is compounded by the preclusion of individuation/singling out as constituting identifiability.[8]

Proposal does not address the purpose of de-identification

The discussion in Section 4.5 of the report does not address the primary use-case for de-identification and therefore fails to consider the appropriateness of allowing de-identified data to exist outside the protections of the Privacy Act – even considering the proposed enforcement of APP 11.1 and APP 8 to de-identified data.

De-identification can be used to avoid consent, and specifically in the context of the Privacy Act, as a loophole to avoid APP 6. By claiming data to be *de-identified* an APP entity may use that data for any purpose without consent nor any other exemptions specified in APP 6.

This fundamentally undermines the autonomy of the individual to limit how their data is to be used, and cements an asymmetric power imbalance between individuals and APP entities in favour of APP entities by codifying as permitted the loophole they are currently exploiting.

It also fundamentally undermines the effectiveness of APP 6 as an APP entity can assert data to be de-identified through a functional process that leaves the bulk of the value of the data and, by extension, its privacy invasiveness intact.

De-identification is not panacea

The process of anonymisation is not about zero risk of re-identification or indeed zero information transfer. It is about ensuring that the information no longer represents a frame of reference at the level of the individual, and by extension, does not require individual consent for its use. Furthermore, there is no widely-accepted definition of *de-identification* that would eliminate the risk of re-identification.

---

[6] See Samarati and Sweeney (n x) 'Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression'.

[7] Ibid.

[8] Section 4.3.4 of the report; See our comments in relation to proposal 4.4.

Effective technologies do exist such as for protecting the integrity of data at rest, in transit, or while being processed by third parties. For example, differential privacy[9] allows the release or sharing of aggregate statistics, and for training statistical or AI models on sensitive data without compromising privacy. What these techniques have in common is a clear definition of privacy properties to be preserved (that can be verified or falsified) that align with an identified threat model.

Reasonable standard of re-identification must consider current best practice

Our group has previously demonstrated re-identifications of 'de-identified' individuals by simply looking up their data from a 'de-identified' dataset in a public dataset. [10] In this case the 'de-identification' could have been framed as a process (satisfying the proposal), that to the data curator represented a defence against reasonable approaches: names were removed, billing dates were shifted. Nonetheless, individuals remained reasonably identifiable.

"Current context" in the proposal implies that, had there been no public information for re-identifying individuals at the time of initial 'de-identified' release, that the release would have been appropriate. It is reasonable to assume that this public data would eventually exist for some of the impacted individuals. "Best available practice" is a welcome addition and should either obviate the reference to "reasonable" identifiability, or at least be included to substantiate what is considered 'reasonable' for purposes of statutory interpretation.

Besides inference of personal information from de-identified datasets, a plethora of works showed that personal data can be inferred from less obvious releases of information – further demonstrating that personal data can be revealed to a malicious third party even when information is revealed in a different form. For example, AI and machine learning technology use personal data to facilitate individuals' experiences such as predictive keyboard. It has been shown that exposure of this technology (API access to a machine learning algorithm) can lead to reverse-engineering it, revealing personal details such email, name, medical information.[11]

Case study: K-anonymity

The current status quo leaves individuals unnecessarily at risk of privacy harm. Consider the continuing use in the field of techniques based on *k*-anonymity, which seeks to render any individual within a dataset indistinguishable from *k*-1 other individuals. The fundamental issues with *k*-anonymity are widely known.[12] However, the biggest issue is not just in the design of *k*-anonymity but in its incorrect application. As a result, many organisations are claiming data is 'de-identified' when it is not and are enjoying a level of data utility that could not be achieved if the data were appropriately protected according to best practice.

To elaborate further on the example of *k*-anonymity (which itself is not regarded by experts as adequate), consider an example where each individual has multiple transactions. If correctly applied,

---

[9] Cynthia Dwork et al, 'Calibrating Noise to Sensitivity in Private Data Analysis' (2006) 3876 *Theory of Cryptography* 265

[10] Culnane, Rubinstein and Teague, (n 3) *Health Data in an Open World*.

[11] Shokri et al, 'Membership Inference Attacks Against Machine Learning Models' [2017] *IEEE Symposium on Security and Privacy* 3; Carlini et al 'Extracting Training Data from Large Language Models' (Conference Paper, 14 December 2020) https://doi.org/10.48550/arXiv.2012.07805.

[12] Chris Culnane and Kobi Leins, 'Misconceptions in Privacy Protection and Regulation' (2019) 36(2) *Law In Context* 49.

each individual would have all their transactions combined into a single row and then that row would have to be indistinguishable from *k*-1 others. This creates a frame of reference at the individual level[13]. The incorrect application of *k*-anonymity is to treat each transaction as a tuple: It is common for transactions to be made indistinguishable, not individuals. In effect this incorrect approach uses individual transactions as the frame of reference and applies the privacy protection only at that level. This tends towards failure as the set of transactions associated with an individual remains discernible and it is that set that individuates the person and allows them to be distinguishable and therefore targeted or profiled. The frame of reference needs to be at the level that the privacy protection is seeking to protect, i.e., it must be at the individual level not the transaction level. In the cases of the Medicare/PBS dataset[14] and Myki dataset[15], our group was able to reidentify individuals based on multiple transactions.

However, as mentioned, *k*-anonymity does not match a specific threat model and preventing singling out of individuals does not inherently protect them from harm (even if used at individual level).

Proposal 4.7 *Consult on introducing a criminal offence for malicious re-identification of de-identified information where there is an intention to harm another or obtain an illegitimate benefit, with appropriate exceptions.*

We strongly recommend against criminalising re-identification, for it is ineffective at protecting privacy for several reasons: such a law requires a definition of *de-identification*; re-identification will not necessarily be detectable; such a law would not apply to entities outside Australian jurisdiction; it would stifle whistle-blowers.

The root cause of privacy failures is poor 'de-identification' and poor data management. It is this behaviour that should be addressed, not subsequent re-identifications. Were a dataset protected with strong guarantees, then re-identification should not be possible to the extent requiring criminalisation.

A law criminalising re-identification will depend on a definitive definition of 'de-identified', since to commit the offence of re-identification, first the dataset must have been considered to have reached some threshold or state of having been de-identified.

Re-identification of de-identified data would only rarely be detected. There are no easily detectable artefacts of re-identification, unless re-identifications are publicised by white-hats or victims of ineffective de-identification looking for compensation. It would be trivially easy to obfuscate where any commercially re-identified data insights came from. Absent an admission, adjudicators would not be able to establish the provenance of information.

If "illegitimate benefit" is to be defined,[16] then legitimate benefits would not be covered. We struggle to understand what these might be. Would "intention to harm" cover an invasion of privacy?

---

[13] Pierangela Samarati and Latanya Sweeney, 'Protecting Privacy when Disclosing Information: k-Anonymity and Its Enforcement through Generalization and Suppression' (Technical Report, Computer Science Laboratory SRI International, 1998)

[14] Culnane, Rubinstein and Teague (n 3) *Health Data in an Open World*

[15] Chris Culnane, Benjamin IP Rubinstein and Vanessa Teague, *Stop the Open Data Bus, We Want to Get Off* (Report, 16 August 2019) <https://doi.org/10.48550/arXiv.1908.05004>.

[16] Proposal 4.7; page 40 of the report.

Section 4.5.2, when discussing the inclusion of an APP Prohibition, states:

> *"While an offence for malicious re-identification would be directed at deterring bad actors, the potential for privacy harms resulting from re-identification also applies where APP entities share de-identified information with third parties who engage in re-identification."*

If a third party engages in re-identification then presumably, they are causing a privacy harm to the individual (unless they are identifying poor de-identification practices of an APP entity).

Were such an offence to be introduced, it is of paramount importance that it does not impinge on reasonable public interest evaluation of the effectiveness of any claimed de-identification method through research. It should also not inhibit the evaluation of re-identification possibilities on any such dataset.

Whilst this is noted in the review as being one of the exceptions, the exceptions themselves are too broad. For instance, we are unclear of the intention behind excluding data analysis. This risks permitting re-identification for the purpose of research on the data subjects, which should obviously not be permitted. At the same time, the currently stated exemption might not permit commercial security consultants to perform audits of de-identification practices.

Proposal 4.8 Prohibit an APP entity from re-identifying de-identified information obtained from a source other than the individual to whom the information relates, with appropriate exceptions. In addition, the prohibition should not apply where: (a) the re-identified information was de-identified by the APP entity itself - in this case, the APP entity should simply comply with the APPs in the ordinary way. (b) the re-identification is conducted by a processor with the authority of an APP entity controller of the information.
This proposal potentially raises concerns about what would be prohibited in Proposal 4.7, and the prohibition and enforcement of any re-identification restrictions.

Firstly, where de-identified data is supplied from the individual, would permitting re-identification undermine APP 2? If an individual has chosen to only provide de-identified data, it would appear to be a breach of their wishes to subsequently re-identify it and undermine their autonomy.

This would appear to exclude from the prohibition a current widespread practice that harms individual privacy, where platforms collect data from users under the guise of it being "de-identified" or "non-identifiable" and subsequently linking it to their identifiable profiles when they log into the platform. No justification is given for this exclusion.

Second, excluding re-identification of data de-identified by the APP itself appears to undermine APP 11.3. If the data had to be destroyed or de-identified after it was no longer needed for the purpose it was collected, in what circumstances would it be permissible to reverse that principle? If 'de-identified' data were no longer about an identifiable individual, then this should not be necessary.

Third, no explanation is offered as to how such data could be treated like other data under the APP. If the data is re-identified what purpose is attached to it? If it is the original purpose, then the APP entity would appear to be in breach of APP 11.3. This creates significant ambiguity about what is going to be permitted.

For example, imagine an individual engages with an APP entity for a specific purpose – Purpose A, and when their data is no longer needed for that purpose the APP entity de-identifies it. The individual subsequently engages the same APP for a different purpose, Purpose B. If the APP entity re-identifies the data from Purpose A and attaches it to the identifiable data they are currently permitted to hold for Purpose B is that data now covered by Purpose B or Purpose A?

If such actions were to be permitted, this would seem to allow APP entities to undertake data enrichment using data that they were only able to hold because they had claimed it was no longer reasonably identifiable. This contradicts the claim that it was no longer reasonably identifiable. The intent of APP 11.3 is that the data should no longer be used in the frame of reference for that individual by the APP. Permitting re-identification of it would fundamentally contradict the spirit of APP 11.3.

Privacy Policies and collection notices

Proposals 10.1, 10.2 and 10.3

Pro-privacy defaults instead of reliance on consent

The report's suggestion of improving the substance and presentation of privacy policies is worthwhile (10.2, 10.3), however an approach that focuses on consent as grounds for any data processing is fundamentally misguided.

As numerous works have noted, the ability of consumers to meaningfully consent to the myriad of uses is hampered by both the feasibility and by the extent to which it is rational for a user to spend their time as such. [17] Users do not read privacy policies or statements, and nor should we expect them to.[18] Further, users have little bargaining power when interacting with APP entities and thus consent is largely illusory.[19]

Rather, it is necessary to introduce mandatory pro-privacy defaults to ensure that consumers are adequately protected. We do not oppose providing users with information as to the products they use– indeed this is vital–but a regime that relies on user consent and opt-outs is one that is already fundamentally compromised.

Research

Proposals 14.1, 14.2, 14.3

We refer you to our comments in respect of **Proposal 4.7**.

A broad exemption should not permissive of re-identification for the purpose of research on the data subjects. At the same time, the exemption as stated in the report might not permit commercial security consultants to perform audits of de-identification practices.

---

[17] Katherine Kemp, '94% of Australians do not read all privacy policies that apply to them – and that's rational behaviour' *The Conversation* (online, 14 May 2018) https://theconversation.com/94-of-australians-do-not-read-all-privacy-policies-that-apply-to-them-and-thats-rational-behaviour-96353, citing CPRC, *Consumer Data and the Digital Economy* (Report, 27 July 2018).

[18] Ibid.

[19] See eg Robert A Hillman and Jeffrey J Rachlinski, 'Standard-Form Contracting in the Electronic Age' (2002) 77(2) *New York University Law Review* 429; David Hoffman, 'Defeating the Empire of Forms' (2023) *Virginia Law Review* (forthcoming).

Security, retention, and destruction

Proposals 21.2, 21.3, 21.5

We are broadly supportive of strengthening the OAIC guidelines, and in particularly enhancing guidance to more clearly articulate *non-exhaustive* baseline security requirements and steps to ensure protection and secure deletion of data (21.2, 21.3, 21.5).

However, we emphasize that any articulated minima must not supplant a requirement to ensure that measures are reasonable within the context in which they are taken, i.e., adhering to the baseline requirements should not on its own constitute a safe harbour.[20]

Measures that are adequate to provide reasonable reduction in risk are sensitive to both the nature of the data processed/stored, and the accompanying technical/organizational infrastructure. What may be sufficient in one context, may be grossly inappropriate under a different threat model.

In relation to deletion of data, we argue that the guidelines include a statement to the effect that generic security concerns are inadequate ground for an entity to reject their obligation to delete data. Even though user data may potentially be of use in anomaly detection or in fraud prevention, the greater risk to individuals stems from unnecessary retention.

Proposal 21.1

The report's proposal to amend APP 11.1 to state that the *reasonable steps* requirement includes technical and organisational measures (21.1) is unlikely to result in meaningful improvements on the ground.

As previous submissions noted, organizations widely appreciate that a mix of governance and technical controls are needed to adequately protect data. While amending the APP is unlikely to result in any negative impacts, the positive impacts too, are likely to be minor.

[20] As a matter of regulatory design, safe harbour provisions stipulate certain requirements an entity must comply with in order to receive the benefit of statutory protections from liability. Provisions within the US Digital Millennium Copyright Act (DMCA) are a prime example.

Direct Right of Action

Proposal 26.1 *Amend the Act to allow for a direct right of action in order to permit individuals to apply to the courts for relief in relation to an interference with privacy. The model should incorporate the appropriate design elements discussed in this chapter.*

We support the introduction of a direct right of action; however, we argue it should ***not*** be subject first to a complaint to the OAIC.

The complaint gateway will erect a burdensome administrative hurdle, which will obviate the stated purpose of the DRA "… to enhance individuals' control of their personal information and reflect current community expectations"[21] by permitting individuals to enforce the Act directly. The propensity for a gateway to become such an inhibition is illustrated by current experience with Freedom of Information requests, whereby there are some 587 unresolved requests older than three years.[22]

The concern for sparing court resources is unfounded in this context.[23] Firstly, the well-known prohibitive costs of litigation will dissuade most potential claimants from making a claim, let alone pursuing it to litigation. The report itself acknowledges this.[24] Frivolous and vexatious claims, should they make it to Court, can thus be adequately guarded against by the proposed requirement to seek leave of court.

Furthermore, should claimants pursue redress though alternative dispute resolution (ADR), the effect of the complaint gateway will weaken their position in these forums also. This is because ADR, while not occurring in a court of law, still takes place under the 'spectre of litigation' – that is, parties negotiate on the basis that their rights may nonetheless be affirmed in a court of law if the matter cannot be settled, especially when one party acts in bad faith.

Should one party know the other is unlikely to pursue their claim to litigation (because of onerous administrative requirements imposed by the complaint gateway), this may confer the power of one party to extract concessions from the other, resulting in undesirable outcomes.

The report again expressly acknowledges this point, in stating that the DRA ought to increase consumer's bargaining power with businesses and encourage greater compliance with the Act.[25] Impediments on recourse to the courts would obviate this purpose.

As far as the proposals seek to strike '…an appropriate balance between improving individuals' access to the courts, discouraging unmeritorious claims and efficient use of court resources', the imposition of a complaint gateway does not achieve this balance. It skews the balance too far in favour of the latter considerations, at the expense of upholding individuals' right to secure compliance with the Act in a court of law, should they see fit to do so.

---

[21] Page 273 of the report.

[22] Christopher Knaus, 'Australia's FOI backlog: 587 cases remain unresolved more than three years on' *The Guardian* (Online, 21 March 2023).

[23] 273: 'While there would be costs associated with a direct right of action … these may be mitigated by ensuring that the design of a right of action strikes an appropriate balance between improving individuals' access to the courts, discouraging unmeritorious claims and efficient use of court resources.

[24] Page 276 of the report, [26.2.4].

[25] Page 273 of the Report under 2.6.