



Algorithm Charter for Aotearoa New Zealand

Year 1 Review

20 December 2021



ACN 087 047 809
ABN 29 087 047 809
www.taylorfry.com.au

Taylor Fry Pty Ltd



Table of contents

1	Executive summary	4
1.1	Charter specifics	4
1.2	Capability and capacity	5
1.3	Engaging with Treaty Partners	6
1.4	Public awareness	6
1.5	Wider context	7
2	Summary of review purpose and methodology	8
2.1	Purpose and scope	8
2.2	Methodology	9
3	Experience of agencies	10
3.1	Agency context	11
3.2	Key review questions	12
3.3	Non-signatory agencies	18
3.4	Interviews with subject matter experts	19
3.5	Document review	20
4	Key themes and considerations	23
4.1	Charter specifics	24
4.2	Capability and capacity	26
4.3	Engaging with Treaty Partners	28
4.4	Public Awareness	29
4.5	Wider Context	30
4.6	Consolidated list of considerations	32
Appendix A	Review Terms of Reference	35
Appendix B	Review Methodology	38

1 Executive summary

The review involved interviewing 25 agencies and subject matter experts, surveying other agencies and reviewing a range of agency and external documentation. It has elicited a large amount of insight and information which we have organised into key themes and considerations in Sections 1.1 to 1.5 of this summary.

More detail can be found in Sections 3 and 4. The review's purpose and methodology are summarised in Section 2.

Overall, the following key messages have been made clear by the review:

- **There is almost universal support for the charter.** Agencies and subject matter experts strongly support the charter in its role of facilitating public awareness and public trust in government's use of algorithms.
- **Agencies have made some progress in implementing charter commitments.** Particularly in terms of performing algorithm stocktakes and categorising algorithms by risk. Given the charter was introduced in July 2020, progress has been reasonable. However, there remains a lot of implementation work to ensure compliance with all charter commitments.
- **Agencies are implementing charter commitments in isolation.** There is little sharing of information and best practice standards leading to some inefficiencies and inconsistencies.
- **Capability and capacity limitations will limit the rate of implementation progress** – The charter introduces a range of complex considerations requiring expert input. Capability and capacity both within agencies and the expert community is limited.

In more detail, the review identified the following 13 themes:

1.1 Charter specifics

Theme 1: Most agencies and subject matter experts interviewed see real value in the charter

The fact that the charter is an across-government initiative helps them to manage public perception and social licence. It also provides a useful framework for discussions between data teams and senior management for resourcing and planning.

Considerations

- Make sure the charter continues to be transparent and public facing
- Continue to update the public on developments with respect to the charter.

Theme 2: There is confusion as to what should be considered an algorithm and captured by the charter

Many agencies have found it challenging to establish what they should and shouldn't capture under the charter.

Considerations

- Maintain the current approach of not providing an explicit algorithm definition. However, supplement existing guidance with further information that clarifies some of the points of confusion and provides more examples of algorithms, tools and processes that should and shouldn't be captured under the charter.

Theme 3: Most agencies view the risk matrix as a relatively simplistic tool with limited use for their context

The risk matrix was a good starting point for agencies thinking about the risk of algorithms. However, as agencies gain maturity in their application of the charter, it might prove too limited.

Considerations

- Consider developing a more detailed risk assessment tool for triaging. Parts of existing assessment frameworks such as the PHRaE and privacy impact assessments, or the Canadian Algorithm Impact Assessment, may be a useful guide.
- Ensure the possible value of algorithms is also captured in risk assessment.

1.2 Capability and capacity

Theme 4: Agencies would welcome a community of practice to support compliance with the charter

One of the strongest themes that came through in the review was agencies' desire for a community of practice for knowledge and information sharing.

Considerations

- Facilitate a community of practice for signatories or potential signatories of the charter. The purpose of the community of practice would be for agencies to share examples, use-cases and problems they have encountered.

Theme 5: Measuring bias and ensuring appropriate human oversight of algorithms is not something in which all agencies have expertise

Measuring bias is a key feature of algorithmic oversight. However, how to measure bias is a complex technical and philosophical problem. Some agencies have struggled to find the expertise to resolve these issues and make trade-offs between different types of bias.

Considerations

- Consider building on the sketch in the “*Algorithm Implementation Guide*”, with a more detailed resource guide for agencies to use. Where appropriate, this could form guidelines and principles for best practice, including technical descriptions of bias assessment protocols.
- A guide to suitable open-source software libraries that provide tools for evaluating bias and interpreting the output of algorithms would be useful.

Theme 6: An oversight body to provide support and oversee aspects of the charter would be seen as beneficial

The Subject Matter Experts (SMEs) and agencies interviewed identified a lack of a clear oversight body for agencies to go to for advice on implementing charter commitments, and for oversight of its principles going forward.

Considerations

- Consider the creation of an oversight body for the algorithm charter. Consider where the function of such a body would naturally sit. For example, does it make sense to sit with the office of Government Chief Data Steward? There is overlap with the function of the Data Ethics Advisory Group, which is another factor in the consideration. An important function of an oversight body would be to liaise with experts in New Zealand and overseas. Many other jurisdictions are working through similar issues to those under consideration in this review.

1.3 Engaging with Treaty Partners

Theme 7: The capacity of experts to support agencies with the partnership commitment and Māori data sovereignty considerations is limited

Both agencies and SMEs identified capacity constraints in seeking experts in Māori data and experts with te ao Māori perspectives. The number of available experts is relatively small and the same people are regularly called upon for advice.

Considerations

- Consider development of best practice in the context of the partnership commitment for sharing among agencies so as to make more efficient use of experts' time and to facilitate consistency
- Clarify the role that existing consultation groups and forums have in helping agencies apply the partnership commitment
- Help agencies to fulfil the principle of parity when engaging in consultation
- Ensure government agencies have realistic expectations about timeframes for consultation with experts, as consultation involves the lengthy process of relationship building.

Theme 8: Many agencies are not clear on how to practically implement the partnership commitment

Agencies have expressed difficulty understanding how to apply the commitment to concrete situations.

Considerations

- Consider working with Māori data experts to develop a more detailed guidance white-paper on what the partnership commitment should consist of, including concrete examples
- A community of practice could circulate and discuss examples of best practice in embedding te ao Māori principles
- Remind agencies that Māori should be involved at the very beginning of the algorithm design process.

1.4 Public awareness

Theme 9: Public reporting of algorithms in use is fragmented and incomplete

Some agencies have taken proactive steps to disclose details about their algorithms on their websites, and a government-wide public algorithm register would build on this work.

Considerations

- Develop an annually updated register of algorithms covered by the charter. The register could be maintained within the Government Chief Data Steward's office.

Theme 10: Public awareness of algorithms in use by government agencies is currently limited

Algorithms are a technical topic, so interest is restricted to a small group of specialists. However, the broader public has interest in the use and purpose of government algorithms.

Considerations

- Investigate novel forms of citizen participation such as citizen assemblies, focus groups etc. for measuring public trust and confidence in government algorithm use. The focus groups used for the Digital Council report on automated decision-making is a useful example.

- Encourage agencies to explicitly refer to the commitments of the charter when dealing with media enquiries – this will boost public awareness of the charter and its principles.

1.5 Wider context

Theme 11: The charter is part of a wider ecosystem of data governance frameworks and policies

The charter exists within a broader ecosystem of data governance and ethics frameworks, such as the Social Wellbeing Agency’s Data Protection and Use Policy (DPUP) and the Privacy Human Rights and Ethics framework (PHRaE). Additionally, individual agencies have developed their own internal frameworks which overlap with many aspects of the charter.

Considerations

- Consider where PHRaE, DPUP, Ngā Tikanga Paihere and other frameworks overlap with the charter commitments. Examples from these wider frameworks can be used to give more detailed implementation guidance for the charter.

Theme 12: The light regulatory nature of the charter places limitations on its ability to offer public assurance and facilitate public trust

There is not necessarily a trade-off between compliance burden (arising from enforcement) and innovation. Some greater enforcement might be necessary to keep social licence, which is necessary for the ongoing development of government analytics and algorithm use. The ultimate goal is to embed considerations from the charter in everyday practice, so there are no longer additional compliance tasks.

Considerations

- Consider ways – such as a publicly available register of algorithms – to evolve the charter from its current light regulatory nature to encourage compliance and best practice.
- Some SMEs suggested non-binding audits of charter compliance. A model here could be the archives NZ audit of public record keeping, which agencies are obliged to engage in on a 5-year cycle.

Theme 13: Algorithms sit within a broader algorithmic system incorporating data sourcing, data use, presentation of algorithm output to users, and decision-making

Algorithm design does not take place in isolation, algorithms are built within data storage systems and are trained on data in such ecosystems. Best practice in data governance is a prerequisite for safe algorithm usage. The presentation of algorithm output is presented to users and how this information is used to support decision making is also critically important.

Considerations

- As the lead for data, Stats NZ is well placed to provide advice to agencies on best practice in data governance. This naturally integrates with work the GCDS delivers, such as “Data Standards” and the “Data Stewardship Framework”.
- Advice on monitoring of algorithm changes during the lifecycle should be explicitly included in any further implementation guide.

2 Summary of review purpose and methodology

2.1 Purpose and scope

Stats NZ launched The Algorithm Charter for Aotearoa New Zealand (“the charter”) in July 2020. The charter is one way that government demonstrates transparency and accountability in the use of data. The charter has 27 signatories to date, encompassing more than half of New Zealand’s government departments, including the most significant users of data.

Written into the charter is a commitment to review its operation after twelve months from launch. Taylor Fry has been asked to carry out the twelve months review.

The objective of the review is to learn from the first year of the charter’s implementation. Issues to be covered include: the experiences of agencies, any early indications of positive impacts or unintended consequences, the support needs of signatories, embedding te ao Māori perspectives and reflecting the principles of the Treaty of Waitangi / te Tiriti o Waitangi, and the relationship between the charter and developments in addressing Māori data sovereignty considerations.

Whilst conducting the review, the following matters were explicitly set out-of-scope:

- The text of the charter
- The voluntary nature of charter
- The remit of the charter being limited to the public sector
- Outcomes and value for money evaluation.

In some areas of this report, we have touched on some of the out-of-scope items where we think they are relevant to the broader scope of the review. More detail on the Terms of Reference for the review is contained in Appendix A.

2.2 Methodology

The methodology comprised three information gathering methods:

1. Face-to-face interviews with treaty partners, data ethics specialists (including Māori data specialists) and key stakeholders including officials from Stats NZ, signatory agencies, and key civil society representatives
2. Questionnaires for other stakeholders
3. Analysis of documents that are publicly available or that can be provided by Stats NZ and other agencies, as well as key literature.

We have assimilated the information from these methods and present our findings in the remainder of this report. Section 3 describes the experience of agencies in the first year since launch, and Section 4 presents the key themes and considerations that arose out of the review. More detail on the review methodology is contained in Appendix B.

3 Experience of agencies

The charter contains six core commitments:

1. **Transparency** – Maintaining transparency by clearly explaining how decisions are informed by algorithms
2. **Partnership** – Embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi
3. **People** – Identifying and actively engaging with people, communities and groups who have an interest in algorithms, and consulting with those impacted by their use
4. **Data** – Making sure data is fit for purpose by understanding its limitations and identifying and managing bias
5. **Privacy, ethics and human rights** – Ensuring that privacy, ethics and human rights are safeguarded by regularly peer reviewing algorithms to assess for unintended consequences and acting on this information
6. **Human oversight** – Retaining human oversight by:
 - Nominating a point of contact for public inquiries about algorithms
 - Providing a channel for challenging or appealing of decisions informed by algorithms
 - Clearly explaining the role of humans in decisions informed by algorithms.

These commitments involve a range of complex considerations for signatory agencies to work through. Most of the current signatories signed up to the charter at launch in July 2020 and hence have had a little over a year to work through these considerations.

The review Terms of Reference (see Appendix A) sets out seven key questions. In this section we discuss the experience of agencies with the charter through the lens of these questions.

We also discuss the views of subject matter experts (SMEs) consulted through the review. During interviews with SMEs we discussed some of the feedback from the interviews with agencies, and sought their opinion on solutions to specific issues and challenges.

Many of the points made are discussed further in a thematic structure in Section 4.

3.1 Agency context

Each agency has a different context and degree of use of algorithms, and so responses to the charter have varied significantly

While the review has yielded a range of common themes, explored in Section 4, it is important to note that each agency's context is different. Consequently, the relevancy of the charter to them, and their response, differs. Key factors differentiating agencies include:

- **Whether agencies serve members of the public directly or not** – For agencies serving members of the public directly, algorithms are sometimes used to inform individual-level decisions e.g. whether person A gets service X or Y.
- **The degree of sophistication of algorithm use** – Some agencies have a long history of sophisticated algorithm use. Others barely use algorithms and/or are just starting to think about how algorithms can add value to their processes.
- **Size and scale** of agencies' operations and resources.

In considering how to comply with the charter commitments, no two agencies are exactly the same. Some degree of tailoring is required.

It is also important to note that the charter is still relatively new having been launched in July 2020. Consequently, it is expected that agencies will be working through how to implement the charter for some time to come.

3.2 Key review questions

Question:	What successes have signatories had in implementing the six charter commitments to date, and how can these be built on over the next three years?
Summary:	There is clear evidence of the charter having positive effects and resulting in meaningful actions. However, most agencies still have a long way to go before becoming fully compliant with charter commitments.

Contemporaneous to the period of the charter (and indeed before the implementation of the charter) there has been a wider conversation within the public sector about data ethics, privacy and accountability. Several frameworks have emerged in recent times, such as:

- The Social Wellbeing Agency’s Data Protection and Use Policy (DPUP) – This policy describes what ‘doing the right thing’ looks like when collecting or using people’s data and information
- The NZ Privacy, Human Rights and Ethics Framework (PHRaE) – An assessment process of the privacy, human rights and ethical impacts of using personal information to develop new services
- Stats NZ’s Ngā Tikanga Paihere – A framework guiding ethical and culturally appropriate data use
- Other internal agency model development governance frameworks.

These frameworks all serve different purposes, but collectively form part of the overall government ecosystem of data management. Furthermore, most agencies operate other data management processes and governance forums, in addition to cross-agency forums, broader sectoral forums, central legislative controls and specific government stewardship roles such as the Government Chief Data Steward.

While the Charter is a standalone document in that it does not directly link to other legislation or regulation, it does not stand alone in its attempt to govern use of algorithms. Since algorithm use and data management are intrinsically linked, the charter is part of this ecosystem.

Many aspects of the ecosystem overlap with the charter. Consequently, many agencies’ processes that existed when the charter was introduced support compliance with the commitments. Indeed, some agencies have chosen to subsume charter commitments into their existing processes, rather than create new ones. It is not always clear whether changes in agencies’ practice since the charter was introduced are due to the influence of the charter, or part of managing the broader ecosystem and would have occurred anyway.

Having said that, most of the signatory agencies have made some practice changes which were at least partly influenced by the charter. These include:

- Development and implementation of risk management policies
- Initiation of ethics committees and review boards
- Development of frameworks for emerging technologies
- Employment of new staff for implementation
- Performing stocktakes and reviews of existing algorithms
- Applying the risk matrix (or similar internal lens) to evaluate the risk of algorithms
- Disclosure of information on agencies’ websites about algorithms in use.

Some agencies reported other high-level effects of signing the charter, including:

- Increased executive visibility. This has given data teams remit to raise concerns related to the charter early in the algorithm commissioning process.
- Approaching interactions with third-party vendors differently as a response to the charter – e.g. asking more questions about interpretability of third-party supplied models.

Example: Recognising the gaps between frameworks such as PHRaE, DPUP and Ngā Tikanga Paihere and practical implementation, the Ministry of Social Development has developed an in-depth guide to help internal data-scientists apply ethical data science principles. This includes explicit guidance on specific data science methods, as well as how these methods integrate within an organisation. As well as this, the MSD guide includes explicit examples for how to manage potential risks and harms using the principles of PHRaE and the charter.

At a more granular level, based on the individual commitments of the charter, agencies have taken the following actions solely or partly in response to the charter.

Transparency

- Some agencies have lists of algorithms published online. In most cases, this is a new practice and wouldn't have happened without the influence of the charter.
- Some agencies have performed stocktakes and risk assessments and published the output.
- Some agencies have plans for further publishing of information in the future.
- Some agencies publish code online on GitHub.
- Some agencies have published methodological reports online in publicly available academic forums – these explicitly referred to the principles of the charter.

Example: The Ministry of Justice has developed a page “Algorithm use in the Ministry of Justice” on their website. This webpage includes a list of operational algorithms as set out in the charter. Each algorithm has information attached on why it is used, what it is used for, who uses the information, and the assessed risk (based on the algorithm charter risk matrix).

Human oversight

Few agencies use extensive automated decision making without any human input whatsoever. However:

- Where automated decision making does occur, it is mainly for positive or opt-in outcomes.
- Some agencies have implemented monitoring functions (“human-on-the-loop” and “human-in-the-loop”) as part of this commitment.
- When faced with the trade-off between a simpler, less accurate model and a more complex model, some agencies have made the decision to implement simpler models. This is to better fulfil the human oversight commitments of the charter.
- Some agencies have a 2-step process of human checks on outputs from algorithms.
- Some agencies have approached their interactions with third-party vendors differently as a response to the charter. For example, more questions are asked about interpretability of third-party models, as well as clear quantifications of bias.

Partnership

Most agencies have existing consultation processes in place, such as internal Māori expert reference groups; many agencies are also developing external partnership relationships. In some cases, the charter has supported further development:

- Some agencies have appointed (or are appointing) strategic Māori data specialists to take the lead in strategic advocacy for Māori data sovereignty
- Some agencies have engaged in consultation with individual iwi on specific matters relating to algorithms.

People

Some agencies have identified and actively engaged with people, communities and groups who might be affected by algorithms.

Data

- Most agencies have implemented data governance frameworks to better understand and manage limitations of data
- Most agencies recognise that algorithms sit within a wider data ecosystem of warehousing, storage etc. and this effects algorithm design.

Example: New Zealand Police have commissioned a stocktake of algorithms which has been proactively released on their website. This was then used to inform an algorithm governance policy for the future, including model life-cycle management. The charter commitments are specifically referenced in this work.

Implementation

While the charter has clearly had some positive impacts, it is important to note that most agencies are still at a relatively early stage in their implementation of the charter (and/or merging commitments with existing processes). Performing algorithm stocktakes and risk assessments has been the most substantive response to the charter. This means that many of the more complex considerations associated with the charter commitments (e.g. identifying and managing bias) are yet to be addressed. Implementing practical solutions for these is likely to be more challenging and resource intensive.

Question:	Is there any evidence of the Algorithm charter stifling innovations?
Summary:	Most agencies see clear benefits to signing the charter.

Most signatory agencies do not believe there is significant compliance burden imposed by the charter. Most agreed that the charter is currently at the “less stringent” end of the regulatory spectrum. No agency made any reference to the charter stifling innovation. This is partly due to the fact that many agencies are just starting out implementing the charter. Many have performed algorithm stocktakes and risk assessments. But few have implemented actions relating to compliance with specific commitments for specific algorithms. This is not unexpected at this stage of the charter implementation. It is likely that most effort and resources will be needed at this next stage of implementing actions for specific algorithms.

This view wasn’t completely universal. There is at least one significant example of an agency not signing up to the charter due to fears about compliance resource requirements. This particular agency is a sophisticated user of algorithms and so is better placed than most to judge the effort and resources required to fully comply with charter commitments. That said, it seemed that this agency may not have interpreted all elements of the charter as intended, and so may have over-estimated the effort required.

Most agencies see clear benefits to signing the charter – especially in public perception – and so feel the benefits far outweigh any compliance costs at this stage. The benefits identified typically focussed on public perception and trust, which would be a result of agencies changing and improving their practice. The fact that the charter has signatories across government was also identified by many agencies as useful.

Question:

To what extent has the consideration of Māori data sovereignty advanced across the data system over the last 12 months, and what are the implications for successful charter implementation?

Summary:

Successful charter implementation is currently limited by agencies' awareness of how to practically incorporate Māori data sovereignty considerations into their data and algorithm use.

Whilst the charter explicitly excludes issues to do with Māori data sovereignty, government agencies have been starting to develop their position on this, independently of the charter. Many agencies have set up internal Māori reference groups, albeit in most cases their focus is much broader than just data sovereignty considerations. Some agencies also said that they were waiting on the outcomes of related streams of work to inform their approach to Māori data sovereignty considerations, such as the co-design of a Māori data governance model, as part of the Stats NZ Mana Ōrite work programme.

Overall, there appears to be a high level of awareness among agencies of the importance of considering Māori data sovereignty. This awareness has developed over a much longer time period than the last 12 months. However, many agencies are unclear what it means in practical terms for their data management and algorithm use. Most agencies recognise that consultation is a key response, but are less clear about practical steps beyond this.

Many agencies also reported that there is a limited number of experts in this field who can provide advice and support, and this was limiting progress to some degree.

Practical actions and solutions for complying with charter commitments, particularly as they relate to the use of data and embedding te ao Māori perspectives, depend on practical applications of Māori data sovereignty. Hence, successful charter implementation is currently limited by agencies' awareness of how to practically incorporate Māori data sovereignty considerations into their data and algorithm use.

Question:

What, if any, indications of improvements in government transparency and accountability are apparent 12 months after the charter was launched, and to what extent are these attributable to the charter?

Summary:

Some improvements are evident and attributable to the charter.

This question overlaps significantly with the question “*What successes have signatories had in implementing the six charter commitments to date, and how can these be built on over the next three years?*” and was largely addressed in that section.

In addition, a small number of agency interviewees questioned whether the charter was a suitable tool for achieving its purpose in its current form. They felt that on its own, the charter did little to raise the public's awareness of government agencies' use of algorithms, and that awareness was a pre-requisite for trust. The sentiment was that the charter was a good start, but needs to evolve as agencies develop their approaches to compliance and increase their level of algorithm use and maturity (notwithstanding the fact that some agencies are already reasonably sophisticated and mature users of algorithms).

As an example, they suggested that agencies disclosing information about their algorithms on their websites was useful, but was likely to have only a small impact. They felt that more proactive engagement-based tools and processes are required to build public awareness and trust.

Question:

What challenges have signatories faced in implementing the six charter commitments to date?

Summary:

Agencies have experienced a range of challenges in implementing the charter and would like targeted support to help them.

Overall, most agencies feel there is a gap between the high-level principles of the charter, and concrete practice for complying with each of the commitments. Supporting frameworks like PHRaE and Ngā Tikanga Paihere fill some of this gap. However, many agencies are not clear on what would constitute compliance with the commitments to ensure they meet the intended purpose of the charter.

Many agencies also expressed a desire for information on what is considered best practice and access to subject matter expertise. Within this context, the role of bodies like the Data Ethics Advisory Group was questioned and whether they should be part of a support network for agencies.

More specifically, with respect to individual charter commitments, agencies have encountered some of the following challenges:

Transparency

- Algorithm documentation is often too complex for the layperson. This means transparency requires that a range of products are released – including “plain English” documentation. Converting documentation into “plain English” can be challenging.
- Some agencies have conflicts because being transparent could indirectly affect the operation and effectiveness of the algorithms e.g. algorithms that support identifying criminal activity.

Human oversight

For many agencies there are technical problems dealing with this commitment. Quantification of bias and fairness can be complex and is still an evolving research field. There is also limited capability for use of interpretable machine-learning. Currently there isn't a clear source of technical expertise to go to for guidance if agencies need assistance.

Partnerships

Most agencies identified that capacity for the relatively small group of Māori data experts and other experts in te ao Māori was limited. That is, there aren't enough experts in this field to meet all of the support needs of agencies in respect of this commitment. While agencies recognised the significance of this commitment, most agencies were unclear on what practical actions should be taken in the context of algorithm use.

Data

- Some agencies are at a relatively early stage of data governance and management maturity. This limits their ability to use algorithms effectively and consistent with charter commitments. In general, these agencies seem to recognise this and have so far taken a conservative approach to algorithm use. Consequently, the potential efficiency and effective decision-making benefits of algorithms are yet to be fully realised.
- Even though Māori data sovereignty considerations are not explicitly referenced in the charter commitments, it is important to note that data governance is intrinsically connected to Māori data sovereignty considerations. It is not possible to treat algorithm design completely isolated from these considerations.

Question:

What challenges and successes have signatories faced in applying the risk matrix?

Summary:

Most agencies view the risk matrix as a relatively simplistic tool with limited use for their context.

Most agencies view the risk matrix as a relatively simplistic tool with limited use for their context. This view varied to some degree, largely depending on the level of sophistication of algorithm use within an agency - sophisticated users of algorithms were less likely to find it useful.

- Some have chosen not to use it, and instead use their own internal risk evaluation tools.
- Some expressed concern that the risk matrix is too easy to “game” because it is self-assessed. They suggested some independence could be introduced into the risk assessment process. In relation to case studies shared amongst agencies, some agencies expressed concern that the risk level associated with a key algorithm had been underestimated by the responsible agency. This shows there are divergent interpretations of the risk matrix.
- Some agencies view the risk matrix as too high-level for many use cases.
- Some suggested that it feels like the risk matrix underestimates risk, perhaps partly because it measures impact in terms of direct impact on people and does not incorporate other risk impacts that may be important to an agency: for example, risks arising from influencing wider social policy.
- Some agencies felt that the risk matrix was included in the charter to partly compensate for the fact that the charter does not explicitly define an ‘algorithm’. What constitutes a high risk ‘algorithm’ could be more clearly defined if there was more guidance on what the charter does and doesn’t intend to capture.

Question:

What support has Stats NZ provided for the implementation of the charter and its wider adoption, and how can this be improved to assist the charter to meet its potential?

Summary:

Agencies would like more support to help them work through compliance with charter commitments.

Most comments here are about support they would like to receive, rather than any support agencies have received. In particular:

- Most agencies are unclear what constitutes best practice in the context of each charter commitment. They would like more information on this to support their implementation and ensure consistency across agencies.
- Many agencies expressed a desire to see more information sharing between agencies. Some reflected that they were doing what they thought was best but were slightly concerned that other agencies were taking different approaches.
- Many agencies reflected that they did not have an obvious avenue for seeking expert support to help them comply with charter commitments. Many are working through it in isolation.

These three points are clearly linked. Agencies would like more support to help them work through charter compliance.

3.3 Non-signatory agencies

At the date of this report there are 27 government agency signatories to the charter. Some public service departments and departmental agencies, non-public service departments, and crown entities have not signed up to the charter. This is expected given that many government agencies do not use algorithms.

As part of the review, we interviewed and surveyed some government agencies who haven't signed up to the charter. In the main, this confirmed the expectation that they didn't make use of algorithms in a meaningful way and so did not feel it was necessary to sign up to the charter. However, there were some exceptions and other pertinent points gleaned from the review process. In particular:

- We interviewed one non-signatory agency who is known to be a sophisticated user of algorithms. This agency understands the importance of the charter and the safe and ethical use of algorithms more broadly, and is considering signing up to the charter in the near future. They have not signed up to the charter so far, principally because they were concerned that they couldn't meet the resource requirements they felt would be necessary to comply with the charter commitments.
- Some agencies consulted with Stats NZ at the time the charter was introduced and agreed that they were low risk with respect to algorithms.
- Some agencies have expressed an intent to sign up to the charter in the future.
- Some agencies said they knew nothing about the charter. It may be that the people who were initially engaged about the charter when it was introduced have since moved on. It may be useful to have an annual process to reengage with non-signatory agencies that are potential algorithm users.
- Some agencies were not operational at the time the charter was introduced. It may be useful to have an annual process to engage with new agencies on the charter.

3.4 Interviews with subject matter experts

As interviews with agencies progressed, several key themes relating to implementation of the charter emerged. These were then discussed with SMEs, in order to elicit comments and suggestions for how to help apply the charter going forward. SMEs' broader views on algorithm use and the charter were also canvassed.

Most SMEs are generally supportive of the purpose of the charter. They believe it is a good high-level starting point to build on, and more work needs to be done to help agencies with implementation. One SME was more critical of the charter, due to the lack of accountability and vagueness in the wording of commitments. Detailed comments and suggestions from SMEs are incorporated in the themes and considerations discussed in Section 4.

Other comments and suggestions

Some of the comments and suggestions from SMEs, although relevant, don't fit into the main themes we present in Section 4. We have included a summary of these below:

- The risk tool currently uses "impact to New Zealanders" as a criterion, but there are some issues with this when considering the full spectrum of algorithm use. For example, the Department of Conservation might use algorithms targeted to impact the natural environment.
- Aside from other issues to do with Māori data sovereignty and the partnership commitment (which are discussed in Section 4), an issue identified by SMEs was that some Māori data scientists don't have enough access to computing power for algorithm training.
- Human oversight by itself is not a panacea, and it is important for agencies to recognise this. A growing body of research, relating to "the control problem" shows that there are issues of complacency and overreliance for humans involved with hybrid decision-making systems.
- At the moment there is very little opportunity for New Zealanders to get individual recourse on decisions made about them that have been informed by an algorithm. For example, in the U.K. there is the mechanism of "judicial review" which has been used by individuals to challenge decisions made using a facial recognition algorithm. In contrast, for a Privacy complaint, an individual can currently go to the Privacy Commissioner, but there isn't anything like that in the case of algorithms.
- It might be useful to include provisions in the charter to completely exclude certain technologies. For example, in Article 4 of the E.U. Regulation Approach for Artificial Intelligence, several technologies are prohibited, such as "AI systems used for indiscriminate surveillance applied in a generalised manner to all natural persons without differentiation."

3.5 Document review

The purpose of the document review is to examine documents relevant to the interviews: these might be examples mentioned by the interviewees, or examples we identified which could serve as comparisons or exemplars. Some of the examples mentioned by interviewees were internal documents detailing data governance frameworks or systems, while some were externally available or from other jurisdictions.

The purpose of the document review was not to do a literature search on all relevant material dealing with ethical use of algorithms in the public sector.

3.5.1 International frameworks

There are a number of high-level frameworks in operation in other jurisdictions, such as the OECD principles on AI, the US Public Policy Council Association for Computing Machinery principles for Algorithmic Transparency and Accountability, and principles from Google, Microsoft, IBM and Facebook. However, to compare with material from the interviews, we focused on a small selection of comparable frameworks in similar jurisdictions:

- The Australian AI ethics principles¹ - a framework developed by the Australian Government Department of Industry, Science, Energy and Resources with applications mainly to the private sector
- The UK data ethics framework² - A set of principles to guide the design of appropriate data use in the public sector, developed by the UK government
- Canadian government responsible use of artificial intelligence³ - A set of principles, directives and impact assessment tools, developed by the Canadian government.

3.5.2 High level principles can be difficult to implement

Each of the international frameworks is similar to the charter in that they provide standards for agencies to apply to their use of algorithms and are typically not intended to be binding and have no enforceable obligations.

For example, the Australian AI ethics principles are listed as:

- Human, societal and environmental wellbeing
- Human-centred values
- Fairness
- Privacy protection and security
- Reliability and safety
- Transparency and explainability
- Contestability
- Accountability.

Whilst the UK data ethics framework lists three principles:

- Transparency

¹ <https://www.industry.gov.au/data-and-publications/australias-artificial-intelligence-ethics-framework/australias-ai-ethics-principles>

² <https://www.gov.uk/government/publications/data-ethics-framework>

³ <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai.html>

- Accountability
- Fairness.

Both of these frameworks then give some guidance for implementation in terms of examples and use-cases, drawn from the private and public sectors. In the case of the Australian principles, some further guidance is available from the Commonwealth Ombudsman “Automated decision-making better practice guide”⁴.

However, in both these jurisdictions the advice given on implementation is still very high level. Detail is missing about how to get from the high-level principles to concrete implementation. SMEs who were familiar with international best practice concur that there is a preponderance of algorithm governance principles across the world, but that no jurisdiction is clearly leading the way in terms of implementation.

One final example is the Canadian “responsible use of artificial intelligence” which also proceeds from high-level principles:

- Understand and measure
- Be transparent
- Provide meaningful explanations
- Be as open as we can
- Provide sufficient training.

However, in addition to the principles, a more detailed tool is provided for implementation - the Canadian Government Algorithmic Impact Assessment Tool – this serves as an example for some of the considerations noted in Section 4.

3.5.3 Internal agency governance documents

We examined a range of internal agency governance documents, including material on the original consultation process for the charter. We note some of the published examples below:

- NZ Police – Safe and Ethical Use of Algorithms Report⁵
- Ministry of Health – Emerging Health Technology: Introductory guidance for safely developing & using Algorithms in Healthcare⁶
- The Privacy, Human Rights and Ethics (PHRaE) Framework⁷.

Information from these internal frameworks was used to inform Section 4, especially with regard to integration of existing frameworks with the charter.

3.5.4 Other public documents supplied or recommended by interviewees

We also examined these public documents provided by interviewees:

- *Algorithmic transparency and accountability* – Transparency International
- *Algorithmic Accountability for the Public Sector* – Ada Lovelace Institute

⁴ <https://www.ombudsman.gov.au/publications/better-practice-guides/automated-decision-guide>

⁵ <https://www.police.govt.nz/sites/default/files/publications/safe-ethical-use-algorithms-report.pdf>

⁶ https://www.health.govt.nz/system/files/documents/pages/introductory_guidance_-_algorithms_v0.4_-_web.pdf

⁷ <https://www.msd.govt.nz/documents/about-msd-and-our-work/work-programmes/initiatives/phrae/phrae-on-a-page.pdf>

- *Te Pou Matakana v. Whānau Tahī Limited* [2021] NZHC 2942 (Ministry of Health / Whānau Ora high court ruling).

3.5.5 Other background material reviewed

We also examined these background documents relating to use of data and artificial intelligence in New Zealand:

- *Government Use of Artificial Intelligence in New Zealand* - New Zealand Law Foundation and Otago University, 2019
- *Principles for the Safe and Effective Use of Data and Analytics* - Privacy Commissioner and Government Chief Data Steward, 2018
- *Trustworthy AI in Aotearoa – AI principles* - AI Forum New Zealand, 2020
- *Towards trustworthy and trusted automated decision-making in Aotearoa* – Digital Council for Aotearoa New Zealand (2020)
- *Data Protection and Use Policy* – Social Wellbeing Agency
- *Ngā Tikanga Paihere* – Stats NZ
- *Māori Perspectives on Trust and Automated Decision-Making* – Te Kotahi Research Institute
- *Principles of Māori Data Sovereignty* – Te Mana Raraunga
- *Submission on the Review of the Draft Algorithm Charter* – Te Mana Raraunga.

As the interviews for this review were concluding, the WAI 2522 claim report was published (“*The report on the comprehensive and progressive agreement for trans-pacific partnership*”). As such, material from this report could not be considered as part of the review. However, the report considers issues relating to governance of Māori data and this is likely to have implications for work supporting the charter in the future.

4 Key themes and considerations

The review process has elicited a large amount of information, opinions and perspectives on the charter and the use of data and algorithms by government agencies. To summarise this material, we have distilled the information into 13 key themes from the review, together with practical considerations for making improvements going forward. We describe these in this section.

The themes are categorised into five areas, albeit many of the themes overlap in more than one area.

- Charter specifics – Themes related to the specific nature and content of the charter
- Capability and capacity – Themes related to the capability and capacity of signatory agencies and the broader analytical community to support meeting charter commitments
- Engaging with Treaty Partners – Themes related to the partnership commitment and Māori data sovereignty
- Public awareness – Themes related to the interface between the public, government agencies and specific algorithms
- Wider context – Themes related to the charter’s place in the broader system of data and algorithm risk management and its role as a regulatory tool.

In Section 4.6, we have provided a consolidated list of considerations and a rating based on their importance and ease of implementation.

4.1 Charter specifics

The charter went through a consultation process with government agencies and some SMEs before implementation in July 2020. Many of the review participants were involved in that consultation process and remain very engaged in matters relating to data and algorithm use. Consequently, the review garnered lots of useful insights and perspectives on the specifics of the charter and how they have operated in practice.

Theme 1: Most agencies and subject matter experts interviewed see real value in the charter

While much of this report focusses on opportunities for improvement, the point should not be lost that the majority of agencies and SMEs interviewed see real value in the charter. They see opportunity to build on the initial charter implementation and work towards the intended charter purpose of improving government transparency and accountability without stifling innovation or causing undue compliance burden.

For agencies, the fact that the charter is an across-government initiative helps them to manage public perception and social licence. It also provides a useful framework for discussions between data teams and senior management for resourcing and planning. Agencies that are interested in progressing their capability and use of more advanced algorithms see the charter as an essential tool to manage risk in this process.

Considerations

- Make sure the charter continues to be transparent and public facing
- Continue to update the public on developments with respect to the charter.

Theme 2: There is some confusion as to what should be considered an algorithm and captured by the charter

The charter does not provide an explicit definition of an algorithm. Rather, it provides some high-level thoughts on what constitutes an algorithm for the purposes of the charter, and a reference to a New Zealand Law Foundation and Otago University report that discusses different types of predictive algorithms. The definition of an algorithm was a robust topic of discussion during the consultation phase on the draft charter.

Decisions about what should be considered an algorithm obviously effect the charter and associated commitments for signatories. There are pros and cons to the existing charter approach of not explicitly defining an algorithm. An explicit definition may result in algorithms being unintentionally excluded. On the other hand, the lack of a definition and/or clear guidance results in some confusion and inconsistency.

Indeed, several agencies have found it challenging to establish what they should and shouldn't capture under the charter. This has probably resulted in some inefficiency and more effort than may have been necessary. Some agencies have hundreds of processes. Trying to establish if each one should be captured as an algorithm under the charter has been relatively onerous for some agencies. Some key areas of confusion include:

- Algorithms that inform policy and funding decisions (as opposed to operational algorithms that interface with the public)
- Efficiency tools that do not inform decision making directly
- Business rules, including those that have been informed by data, and whether an individual business rule or a collection of business rules constitutes an algorithm
- Other business processes that inform decision-making.

For some agencies, there are blurred lines between algorithm risk management and broader business process risk management. A business process may not be an algorithm but may still inform decision-making and be based on data and/or expert opinion. Arguably the principles of the charter could apply to such processes, but might create an onerous burden on some agencies.

Overall, we believe the benefits of the current approach outweigh the negatives, but needs to be supplemented with some additional guidance.

Considerations

- Maintain the current approach of not providing an explicit algorithm definition. However, supplement existing guidance with further information that clarifies some of the points of confusion noted above and provides more examples of algorithms, tools and processes that should and shouldn't be captured under the charter.

Theme 3: Most agencies view the risk matrix as a relatively simplistic tool with limited use for their context

The charter provides a two-dimensional risk matrix with dimensions for:

- The likelihood of an unintended adverse outcome
- The level of impact if an adverse outcome occurs.

Where an algorithm is judged by an agency to sit on these dimensions determines whether the risk rating for that algorithm is high, medium or low. The intent is for agencies to focus charter compliance on high-risk algorithms first.

The risk matrix was a good starting point for agencies thinking about the risk of algorithms. However, as agencies gain maturity in their application of the charter, it might prove too limited. Many agencies and SMEs felt that a more detailed tool would be useful at the outset of the triaging process. Such a tool could incorporate more detailed questions about the purpose, benefits, risks and impacts of the proposed algorithm. For example, if a user states that the algorithm in question is low risk, then they have to specify why this is so.

Some SMEs pointed out that, currently, Privacy Impact Assessments are useful sources of information for external parties to understand what government systems are in use. A more detailed algorithm risk evaluation tool could function as an “Algorithm Impact Assessment”. The Canadian government has already developed a prototype of such a tool, which could serve as a template⁸.

This subject needs to be considered in conjunction with the previous section about the definition of an algorithm. With a broad capture of algorithms, tools and processes under the charter, performing an “Algorithm Impact Assessment” for every example may be too onerous for some agencies.

Note that some agencies have already developed their own risk assessment tools which are used instead of the risk-matrix.

Considerations

- Consider developing a more detailed risk assessment tool for triaging. Parts of existing assessment frameworks such as the PHRaE and privacy impact assessments, or the Canadian Algorithm Impact Assessment, may be a useful guide.
- Ensure that possible value of algorithms is also captured in risk assessment.

⁸ <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html>

4.2 Capability and capacity

Developing, using and maintaining algorithms in a safe and ethical way involves specialist skills and capabilities. The capability and capacity of agencies and the broader network of analytical resources and experts impacts their ability to comply with charter commitments. Several review participants raised concerns about the supply of resources with these skills and capabilities.

Theme 4: Agencies would welcome a community of practice to support compliance with the charter

One of the strongest themes that came through the review was agencies' desire for a community of practice for knowledge and information sharing. Most agencies have addressed their charter commitments largely on their own and without knowledge of how other agencies were going about it. Several agencies expressed concern that perhaps their approach wasn't optimal and that other agencies may have a better approach. Agencies also expressed a need for information on what is considered 'best practice' in respect to different charter commitments and a medium for accessing expert support when needed.

Collectively, the status quo has created some inefficiency and inconsistency in the process. While each agencies' circumstances and use of algorithms is different, they felt that some 'best practice' information and/or principles would help give them confidence in their approach and support efficient and consistent approaches to meeting charter commitments.

While grappling with the complex issues raised by the charter, concrete examples often prove useful guides. For example, many agencies identified the examples in the Stats NZ document "*Algorithm Charter – Implementation Guidance*" as useful. All agencies and SMEs interviewed agree that sharing examples of best practice, as well as possible issues and solutions would be useful for all signatories of the charter. Several agencies identified concerns about consistent measurement of risk across the sector. A community of practice would likely increase consistency of measurement of risk and implementation of the charter more broadly.

Considerations

- Facilitate a community of practice for signatories or potential signatories of the charter. The purpose of the community of practice would be for agencies to share examples, use-cases and problems they have encountered.

Theme 5: Measuring bias and ensuring appropriate human oversight of algorithms is not something in which all agencies have expertise

Measuring bias is a key feature of algorithmic oversight. However, how to measure bias is a complex technical and philosophical problem. Measurement is also a necessary first stage before making decisions about how much bias is permissible. Some agencies have struggled to find the expertise to resolve these issues and make trade-offs between different types of bias. In a public context, it is also a challenging discussion for agencies to lead and acknowledge that pure human decision-making (which algorithmic systems often replace) are inherently biased.

Similarly, the exact nature of how to ensure appropriate human oversight of algorithms is both a complex technical and philosophical issue. Several of the SMEs interviewed highlighted that incorporating human oversight in an algorithmic system is not fool proof and requires careful thought.

Relatedly, most agencies have capability gaps for critically evaluating solutions that may support bias management, transparency and effective human oversight, such as experience with interpretable machine

learning and fairness measurement. One SME also believed that understanding of broader “data justice”⁹ issues amongst agencies was not sufficiently developed to implement the commitments of the charter.

Considerations

- Consider building on the sketch in the “*Algorithm Implementation Guide*”, with a more detailed resource guide for agencies to use. Where appropriate, this could form guidelines and principles for best practice, including technical descriptions of bias assessment protocols.
- A guide to suitable open-source software libraries that provide tools for evaluating bias and interpreting the output of algorithms would be useful.

Theme 6: An oversight body to provide support and oversee aspects of the charter would be seen as beneficial

While Stats NZ developed and introduced the charter, their role does not formally extend to oversight of charter compliance and the use of algorithms by government agencies.

SMEs and agencies identified a lack of a clear oversight body for agencies to go to for advice on implementing charter commitments, and for oversight of its principles going forward. In the past, the Data Ethics Advisory Group has played some of the roles that would be expected of such a body¹⁰, but there is lack of clarity about whether that is the appropriate forum going forward and the role the GCDS should play in this regard (if any).

Without a clear oversight body, agencies have largely interpreted the charter commitments themselves, creating some inefficiency and inconsistency.

An oversight body would also be the natural forum for receiving any complaints about misapplication of the charter principles. Relatedly, it was noted by SMEs that there are few existing processes for the public to question the application of an algorithm in a decision that directly affects them.

Considerations

- Consider the creation of an oversight body for the algorithm charter. Consider where the function of such a body would naturally sit. For example, does it make sense to sit with the office of Government Chief Data Steward? There is overlap with the function of the Data Ethics Advisory Group, which is another factor in the consideration. An important function of such an oversight body would be to liaise with experts in New Zealand and overseas. Many other jurisdictions are working through similar issues to those under consideration in this review.

⁹ “Data justice” is a framework for engaging with the intersection of datafication and society in a way that privileges an explicit concern for social justice, for more details see <https://policyreview.info/articles/analysis/data-justice> or <https://datajusticelab.org/about/>

¹⁰ For example, the DEAG provided advice to a variety of agencies on the suitability of algorithms <https://www.data.govt.nz/leadership/advisory-governance/data-ethics-advisory-group/guidance-from-data-ethics-advisory-group/>

4.3 Engaging with Treaty Partners

Whilst the charter explicitly excludes important considerations, such as Māori Data Sovereignty, the partnership commitment explicitly references the principles of the Treaty of Waitangi. Most agencies already have internal and external bodies to engage with Māori and fulfill the Crown's Treaty commitments. However, the requirements of the partnership commitment in the charter have presented new challenges for agencies.

Theme 7: The capacity of experts to support agencies with the partnership commitment and Māori data sovereignty considerations is limited

Both agencies and SMEs identified capacity constraints in seeking experts in Māori data and te ao Māori. The number of available experts is relatively small and the same people are regularly called upon for advice and input, often on an unpaid basis.

This issue is exacerbated by the consistency issue referenced elsewhere in this section. Because agencies are largely managing their charter commitments (and broader data governance processes) independently, each agency has their own engagement channel with experts. This likely duplicates work and increases the call on experts' time.

One SME felt that a lack of understanding of Māori data sovereignty principles on the part of government agencies hampered the ability of experts to support agencies. This is related to some of the capability issues identified in Theme 5.

Considerations

- Consider development of best practice in the context of the partnership commitment for sharing among agencies so as to make more efficient use of experts' time and to facilitate consistency
- Clarify the role that existing consultation groups and forums have in helping agencies apply the partnership commitment
- Help agencies to fulfil the principle of parity when engaging in consultation
- Ensure government agencies have realistic expectations about timeframes for consultation with experts, as consultation involves the lengthy process of relationship building.

Theme 8: Many agencies are not clear on how to practically implement the partnership commitment

Both agencies and SMEs have identified issues with the partnership commitment: there is difficulty understanding how to apply the commitment in practice. Agencies are unclear about how well they comply with this specific charter commitment and what actions they need to take to work towards compliance.

Given these concerns – and given the importance of the partnership commitment – it is important that more detail is given to practitioners about how the partnership commitment could be realised. In consultation with Māori data experts, some examples of best practice could be drawn up and disseminated¹¹. International frameworks such as the United Nations Declaration of the Right of Indigenous Peoples could also inform this work. A more detailed overview of what is intended is necessary.

¹¹ Te Mana Raraunga has utilised Māori Data Sovereignty principles to propose recommendations for implementing algorithm/models; these are collected at the end of the Te Mana Raraunga submission to the draft of the charter, and could help inform more detailed guidance.

Considerations

- Consider working with Māori data experts to develop a more detailed guidance white-paper on what the partnership commitment should consist of, including concrete examples
- A community of practice could circulate and discuss examples of best practice in embedding te ao Māori principles
- Remind agencies that Māori should be involved at the very beginning of the algorithm design process.

4.4 Public Awareness

Transparency is a key principle of the charter, which helps facilitate public awareness, a pre-requisite for public trust and social license. There are many dimensions to public awareness, and it is important that this extends beyond the realm of experts and those who frequently advise government agencies.

Theme 9: Public reporting of algorithms in use is fragmented and incomplete

As part of the 2018 algorithm assessment, many agencies engaged in a stocktake of algorithms in use. However, some agencies are still only starting this process, or have decided to focus on a select few examples, rather than a full stocktake of algorithms.

As maturity in this area develops, government agencies will develop tools to record and track algorithms in use and which ones are covered by the charter. An important part of the transparency commitment is for the public to have an accessible record of which algorithms currently in use. Many SMEs also expressed a preference for an explicit public register to ensure easy access and record consistency.

Some agencies have taken proactive steps to disclose details about their algorithms on their websites, and a government-wide public algorithm register would build on this work.

Considerations

- Develop an annually updated register of algorithms covered by the charter. The register could be maintained within the Government Chief Data Steward's office.

Theme 10: Public awareness of algorithms in use by government agencies is currently limited

The charter has spurred greater levels of transparency - many agencies published more information about their use of algorithms online. However, some SMEs and agencies would like to see more awareness among the public of government algorithm use. They believe greater awareness is necessary for public trust and social licence.

Algorithms are a technical topic, so interest is restricted to a small group of specialists. However, the broader public has interest in the use and purpose of government algorithms. Some SMEs suggested these should be openly discussed in public forums and could involve novel forms of consultation and participation.

Considerations

- Investigate novel forms of citizen participation such as citizen assemblies, focus groups etc. for measuring public trust and confidence in government algorithm use. The focus groups used for the Digital Council report on automated decision-making is a useful example¹².
- Encourage agencies to explicitly refer to the commitments of the charter when dealing with media enquiries – this will boost public awareness of the charter and its principles.

4.5 Wider Context

The charter exists within a wider context. There are evolving public discussions about data ethics and algorithm use, as well as broader issues of data governance in the public sector. Positioning the charter within this context will help agencies implement the charter and will help build public trust.

Theme 11: The charter is part of a wider ecosystem of data governance frameworks and policies

The charter exists within a broader ecosystem of data governance and ethics frameworks, such as the Social Wellbeing Agency’s Data Protection and Use Policy (DPUP) and the Privacy Human Rights and Ethics framework (PHRaE). Additionally, individual agencies have developed their own internal frameworks which overlap with many aspects of the charter.

This overlap is not necessarily an issue – the concerns of the charter are part of wider public discussions about the use of algorithms. It is natural that other government initiatives should play a role. The key point is to identify synergies between these various initiatives and the commitments of the charter. Where possible, best practice data governance frameworks should be disseminated among agencies to avoid doubling-up of work and to ensure consistency. Links between supporting and complimentary frameworks should be clearly described to help agencies navigate the ecosystem.

Considerations

- Consider where PHRaE, DPUP, Ngā Tikanga Paihere and other frameworks overlap with the charter commitments. Examples from these wider frameworks can be used to provide more detailed implementation guidance for the charter.

Theme 12: The light regulatory nature of the charter places limitations on its ability to offer public assurance and facilitate public trust

Any regulatory framework sits on a scale: with “high-level, lightly enforced” at one end and “detailed, strictly enforced” at the other. The majority of agencies and SMEs believe that the charter currently sits at the “high-level” end of this spectrum. For example, very few feel there is a significant compliance burden currently associated with implementing the charter. Most agencies also believe that there is an opportunity for the charter to move slowly in the direction of more detailed enforcement. This should be possible without introducing undue compliance burden as agencies develop their approaches. There is not necessarily a trade-off between compliance burden (arising from enforcement) and innovation. Some greater enforcement might be necessary to keep social licence, which is necessary for the ongoing development of government analytics and algorithm use. The ultimate goal is to embed considerations from the charter in everyday practice, so there are no longer additional compliance tasks.

¹² Digital Council for Aotearoa New Zealand. (2020). *Towards trustworthy and trusted automated decision-making in Aotearoa*.

Greater accountability will also provide motivation for experts who may be involved with giving advice on the implementation of the charter – without any accountability it may be hard for experts to see the possible results of engaging in consultation.

Considerations

- Consider ways – such as a publicly available register of algorithms – to evolve the charter from its current light regulatory nature to encourage compliance and best practice.
- Some SMEs suggested non-binding audits of charter compliance. A model here could be the archives NZ audit of public record keeping, which agencies are obliged to engage in on a 5-year cycle.

Theme 13: Algorithms sit within a broader algorithmic system incorporating data sourcing, data use, presentation of algorithm output to users, and decision-making

Algorithm design does not take place in isolation, algorithms are built within data storage systems and data warehousing frameworks. The all-important training data for fitting algorithms is built in such systems. Most importantly, data warehousing systems require continual monitoring and lifecycle maintenance: new data is constantly being added, data schemas are updated, and counting rules documented.

As such, best practice in data governance and management of data is a prerequisite for safe algorithm usage. Governance frameworks need to be in place to ensure that agencies have a good view of how and why their data is being used and for what purpose, and to monitor changes over time. This is also a prerequisite for efficient use and codification of business rules.

The presentation of algorithm output to users and how this information is used to support decision making is also critically important. For example, is the output presented with enough information for the user to be able to judge that the data inputs look correct?

Considerations

- As the lead for data, Stats NZ is well placed to provide advice to agencies on best practice in data governance. This naturally integrates with work the GCDS delivers, such as “Data Standards” and the “Data Stewardship Framework”.
- Advice on monitoring of algorithm changes during the lifecycle should be explicitly included in any further implementation guide.

4.6 Consolidated list of considerations

We provide a consolidated list of the considerations made in Table 4.1. To help with prioritisation, we have rated each consideration on two factors:

- How important we believe the consideration is on a five-point scale where 5 is the highest rating.
- The ease of implementing the consideration, where ‘ease’ reflects both the effort required and how clear a solution is. Five-point scale where 5 is the easiest rating.

The ratings are necessarily subjective, though hopefully will help prioritise your resources given that actioning every consideration at once is not possible.

Table 4.1 – Consolidated list of recommendations

Theme	Consideration	Importance	Ease	Page ref.
1	Make sure the charter continues to be transparent and public facing.	5	4	24
1	Continue to update the public on developments with respect to the charter.	4	5	24
2	Maintain the current approach of not providing an explicit algorithm definition. However, supplement existing guidance with further information that clarifies some of the points of confusion noted above and provides more examples of algorithms, tools and processes that should and shouldn’t be captured under the charter.	4	3	24
3	Consider developing a more detailed risk assessment tool for triaging. Parts of existing assessment frameworks such as the PHRaE and privacy impact assessments, or the Canadian Algorithm Impact Assessment, may be a useful guide.	3	3	25
3	Ensure the possible value of algorithms is also captured in risk assessment.	2	3	25
4	Facilitate a community of practice for signatories or potential signatories of the charter. The purpose of the community of practice would be for agencies to share examples, use-cases and problems they have encountered.	5	4	26
5	Consider building on the sketch in the “ <i>Algorithm Implementation Guide</i> ”, with a more detailed resource guide for agencies to use. Where appropriate, this could form guidelines and principles for best practice, including technical descriptions of bias assessment protocols.	5	2	26
5	A guide to suitable open-source software libraries that provide tools for evaluating bias and interpreting the output of algorithms would be useful.	2	5	26

Theme	Consideration	Importance	Ease	Page ref.
6	Consider the creation of an oversight body for the algorithm charter. Consider where the function of such a body would naturally sit. For example, does it make sense to sit with the office of Government Chief Data Steward? There is overlap with the function of the Data Ethics Advisory Group, which is another factor in the consideration. An important function as such an oversight body would be to liaise with experts in New Zealand and overseas. Many other jurisdictions are working through similar issues to those under consideration in this review.	4	1	27
7	Consider development of best practice in the context of the partnership commitment for sharing among agencies so as to make more efficient use of experts' time and to facilitate consistency.	4	2	28
7	Clarify the role that existing consultation groups and forums have in helping agencies apply the partnership commitment.	3	3	28
7	Help agencies to fulfil the principle of parity when engaging in consultation.	3	3	28
7	Ensure government agencies have realistic expectations about timeframes for consultation with experts, as consultation involves the lengthy process of relationship building. Consultees are often less resourced, compared to government agencies.	4	3	28
8	Consider working with Māori data experts to develop a more detailed guidance white-paper on what the partnership commitment should consist of, including concrete examples.	4	2	28
8	A community of practice could circulate and discuss examples of best practice in embedding te ao Māori principles.	5	3	28
8	Remind agencies that Māori should be involved at the very beginning of the algorithm design process.	4	5	28
9	Develop an annually updated register of algorithms covered by the charter. The register could be maintained within the Government Chief Data Steward's office.	4	2	29
10	Investigate novel forms of citizen participation such as citizen assemblies, focus groups etc. for measuring public trust and confidence in government algorithm use. The focus groups used for the Digital Council report on automated decision-making is a useful example.	4	2	29

Theme	Consideration	Importance	Ease	Page ref.
10	Encourage agencies to explicitly refer to the commitments of the charter when dealing with media enquiries – this will boost public awareness of the charter and its principles.	3	3	29
11	Consider where PHRaE, DPUP, Ngā Tikanga Paihere and other frameworks overlap with the charter commitments. Examples from these wider frameworks can be used to provide more detailed implementation guidance for the charter.	3	2	30
12	Consider ways – such as a publicly available register of algorithms – to evolve the charter from its current light regulatory nature to encourage compliance and best practice.	5	2	30
12	Some SMEs suggested non-binding audits of charter compliance. A model here could be the archives NZ audit of public record keeping, which agencies are obliged to engage in on a 5-year cycle.	3	1	30
13	As the lead for data, Stats NZ is well placed to provide advice to agencies on best practice in data governance. This naturally integrates with work the GCDS delivers, such as “Data Standards” and the “Data Stewardship Framework”.	3	3	31
13	Advice on monitoring of algorithm changes during the lifecycle should be explicitly included in any further implementation guide.	5	3	31

Appendix A Review Terms of Reference

A.1 Introduction

Stats NZ is commissioning an independent review of the first year of the Algorithm Charter for Aotearoa New Zealand (the charter). The resulting report will be made publicly available.

Stats NZ released the charter in July 2020. The charter is one way government demonstrates transparency and accountability in the use of data. The charter, which is voluntary, currently has 27 signatories across the New Zealand government. Signatories are committed to applying the principles of the charter in their work.

More information on the charter can be found [here](#).

A.2 Rationale for the review

The charter makes a commitment to undertake a review 12 months after coming into force to ensure it is achieving its intended purpose of improving government transparency and accountability without stifling innovation or causing undue compliance burden.

The implementation and interpretation of the charter is an ongoing process that needs to respond to emerging technologies and be fit-for-purpose for government agencies. A review provides early information on whether the charter is on-track to achieve its purpose; and areas where implementation can be improved, or further implementation support is required. This includes work agencies may have done around embedding te ao Māori perspectives, reflecting the principles of the Treaty of Waitangi / te Tiriti o Waitangi, and Māori data sovereignty considerations.

A.3 Background to the charter

In 2018, Stats NZ released an Algorithm Assessment Report. Development of the charter followed. It was launched in July 2020, and has 26 signatories to date, encompassing more than half of New Zealand's government departments including the most significant users of data.

The charter contains six commitments that agencies sign up to, covering:

- Transparency
- Partnership (embedding a te ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi / te Tiriti o Waitangi)
- People
- Data
- Privacy, ethics, human rights, and
- Human oversight.

The charter represents a normative approach to establishing best practice across-government. It offers agencies a framework and guidance to help them meet transparency and accountability objectives. As a voluntary, self-regulating agreement, the charter is not legally binding and has no enforcement mechanisms or sanctions. There is also no formal process for monitoring signatory agencies' compliance.

The charter recognises that each signatory is unique. It provides a risk matrix for agencies to identify and focus on algorithms that have a high or critical risk of unintended harms for New Zealanders. On signing the charter, signatories are provided with implementation guidelines.

The charter has been welcomed in New Zealand and internationally. Nevertheless, there are challenges with operationalising it. As a world-leader in this area, New Zealand has not had the benefit of others'

experiences in developing and implementing similar charters. Therefore, it is important to take a structured approach to assessing and learning from the first year of the charter's implementation and identifying opportunities for improvement.

A.4 Review objective

The objective of the review is to learn from the first year of the charter's implementation. This includes: the experiences of agencies, any early indications of positive impacts or unintended consequences, the support needs of signatories, embedding te ao Māori perspectives and reflecting the principles of the Treaty of Waitangi / te Tiriti o Waitangi, and the relationship between the charter and developments in addressing Māori data sovereignty considerations.

A.5 Review questions

The review will examine the following questions:

1. What, if any, indications of improvements in government transparency and accountability are apparent 12 months after the charter was launched, and to what extent are these attributable to the charter?
2. Is there any evidence of the Algorithm charter stifling innovation or causing an undue compliance burden? If so, what?
3. The charter notes that it cannot fully address important complex considerations, such as Māori data sovereignty, and that these require separate consideration. To what extent has the consideration of Māori data sovereignty advanced across the data system over the last 12 months, and what are the implications for successful charter implementation?
4. What successes have signatories had in implementing the six charter Commitments to date, and how can these be built on over the next three years?
5. What challenges have signatories faced in implementing the six charter Commitments to date, and how can these be overcome?
6. What challenges and successes have signatories faced in applying the risk matrix?
7. What support has Stats NZ provided for the implementation of the charter and its wider adoption, and how can this be improved to assist the charter to meet its potential?

A.6 Method

The review will be informed by relevant information sources, including:

- analysis of documents that are publicly available or that can be provided by Stats NZ and other agencies, as well as key literature
- interviews with ○ officials from Stats NZ, signatory agencies and selected non-signatory agencies
- data ethics specialists including Māori data specialists
- key civil society representatives, such as Transparency International New Zealand.

A.7 Scope

The scope of the review is the first 12 months of the charter's implementation and lessons that can be learnt for the next three years. The following matters are out of scope:

- amendments to the text of the charter

- the voluntary nature of charter
- extending the charter to the private sector
- outcomes and value for money evaluation – the charter has not been in place long enough to allow such assessments to be made.

A.8 Reviewer Profile

The review will be undertaken by a suitably qualified independent reviewer appointed by Stats NZ following a procurement process run in accordance with government guidelines.

A.9 Timeframe

The review will commence after the charter has been in place for one year (after 31 July 2021) and will be completed before the end of November 2021.

Appendix B Review Methodology

B.1 Overview

The review scope extends to using all available resources to answer the questions set out in the terms of reference. This involves:

- Face-to-face interviews with treaty partners, data ethics specialists (including Māori data specialists) and key stakeholders including officials from Stats NZ, signatory agencies, and key civil society representatives, such as Transparency International New Zealand
- Questionnaires for other stakeholders
- Analysis of documents that are publicly available or that can be provided by Stats NZ and other agencies, as well as key literature.

B.2 Interviews

The purpose of the interviews is to elicit detailed responses from:

- Signatory agencies, about their experiences with the charter and the set of questions set out in the terms of reference
- Subject matter experts, about their perspectives on the function of the charter after one year.

The interviews compliment other relevant information (such as questionnaires and a review of key documents). In particular, interviews allow us to:

- Canvass opinions and experiences in depth
- Retain flexibility to adjust questions on-the-fly based on circumstances and earlier participant responses.

Table B.1 describes a range of elements relating to the format of the interviews.

Table B.1– Interview format

Item	Approach
Style	The interviews were relatively informal and led by the interviewer.
Location	Where possible, we performed interviews in person. However, depending on COVID-19 alert levels and participants' preferences we performed some interviews by video conference.
Length	Interviews lasted about one hour.
Structure	The interviews started with some standard disclosures. The interviewer then led the participant through a series of question areas. Initial questions in each area were scripted, though the interviewer often asked further unscripted questions depending on participants answers.
Interviewers	Dan Stoner and Callum Sleight of Taylor Fry.
Quality control	Dan and Callum attended the first four interviews together, with Dan leading two and Callum leading two. They then discussed their approaches to ensure consistency from that point forward. Only one of Dan or Callum attended the remaining interviews.

Item	Approach
Recording of interviews	Interviews were recorded so that interviewers could focus on the conversation rather than taking notes. Participants were given the option to not have their interviews recorded if they preferred. Interview audio was only used by Taylor Fry and was not shared with Stats NZ or any other party. The audio files will be deleted at the end of the project.
Disclosures	At the beginning of the interviews a set of scripted disclosures and information for participants were read out. This was to ensure all participants understood the purpose of the review, the purpose of the interview, what was to be covered, how their information will be used and data stored, and choices they have with respect to their data and identity.

The following fourteen agencies were interviewed:

- Te Ara Poutama Aotearoa - The Department of Corrections
- Te Tāhuhu o Te Mātauranga - The Ministry of Education
- Te Tari Taake - Inland Revenue Department
- Te Tāhū o te Ture - The Ministry of Justice
- Te Manatū Whakahiato Ora - The Ministry of Social Development
- Te Tatauranga Aotearoa – Statistics New Zealand
- Toi Hau Tāngata – Social Wellbeing Agency
- Te Kaporeihana Āwhina Hunga Whara – Accident Compensation Corporation
- Hikina Whakatutuki – The Ministry of Business, Innovation and Employment
- Manatū Hauora – The Ministry of Health
- Te Tari Taiwhenua – Department of Internal Affairs
- Te Mana Ārai o Aotearoa – New Zealand Customs Service
- Nga Pirihimana O Aotearoa – New Zealand Police.
- Waka Kotahi – New Zealand Transport Agency

There were two interviews for *Hikina Whakatutuki – The Ministry of Business, Innovation and Employment*, one covering Immigration New Zealand and one covering other business units.

These agencies have been selected for interview on the basis that they are most likely to have high risk algorithms.

The following subject matter experts were interviewed:

- Andrew Sporle, iNZight Analytics Ltd.
- Transparency International New Zealand
- Nessa Lynch, Victoria University
- Colin Gavaghan, Otago University
- Andrew Chen, University of Auckland
- Caleb Moses, Dragonfly Data Science
- Daniel Wilson, University of Auckland

- Māui Hudson, University of Waikato
- Moka Apiti, Digital Navigators Ltd.

There was also an additional subject matter expert who was interviewed but didn't want to be listed in the final report.

B.3 Questionnaires

The purpose of the questionnaires was to reach a broader set of participants than can be practically managed through interviews alone. They don't offer the same level of question flexibility as interviews and tend not to generate as in-depth responses. However, they were a useful medium to capture a broad range of opinions and can be designed to deliver responses that can be analysed as data fields.

There were three questionnaire variations catering to three specific audiences:

1. Analytics Deputy Chief Executives (or equivalent DCE role)
2. Non-signatory agencies
3. Signatory agencies not interviewed

Questionnaire format

Table B.2 describes a range of elements relating to the format of the questionnaires.

Table B.2 – Questionnaire format

Item	Approach
Format	Online delivery
Length	5 to 15 minutes to complete depending on participants' level of response to free text fields.
Structure	Most questions structured with categorical response options. Free text fields used in some cases.
Quality control	Initially, a small number of participants for each questionnaire were invited to respond. These participants were then asked for feedback on the questionnaire. Where possible, improvements were made before inviting the broader set of participants to respond.

Participants

Intended participants for the four questionnaire variations are as shown in the following sections.

Analytics DCE

All DCEs who are part of the analytics DCE group hosted by Stats NZ except those who were interviewed. This incorporates representatives from:

- Te Tāhuhu o te Mātauranga - The Ministry of Education
- Te Tari Taake - Inland Revenue Department
- Te Tāhū o te Ture - The Ministry of Justice
- Oranga Tamariki – The Ministry for Children
- Te Manatū Whakahiato Ora - The Ministry of Social Development
- Tatauranga Aotearoa – Statistics New Zealand

- Te Kaporeihana Āwhina Hunga Whara – Accident Compensation Corporation
- Hikina Whakatutuki – The Ministry of Business, Innovation and Employment
- Te Puni Kōkiri – The Ministry of Māori Development
- Te Tari o te Pirimia me te Komiti Matua - Department of Prime Minister & Cabinet
- Te Manatū Mō Te Taiao – The Ministry for the Environment

Non-signatory agencies

Selected agencies who are not currently signatory agencies:

- Te Tai Ōhanga - The Treasury
- Te Papa Atawhai - Department of Conservation
- Te Tira Tiaki - Government Communications Security Bureau
- Manatū Taonga - Ministry for Culture and Heritage
- Manatū Ahu Matua - Ministry for Primary Industries
- Manatū Kaupapa Waonga - Ministry of Defence
- Te Pā Whakamarumaruru - New Zealand Security Intelligence Service
- Te Kawa Mataaho - Public Service Commission
- Te Tari Hara Tāware - Serious Fraud Office
- Te Aho o te Kahu - Cancer Control Agency
- Te Rākau Whakamarumaruru - National Emergency Management Agency
- Te Tari Mātāwaka - Ministry for Ethnic Communities
- Callaghan Innovation
- Te Mana Rererangi Tūmatanui o Aotearoa - Civil Aviation Authority of New Zealand
- Earthquake Commission
- Mana Tohu Matauranga o Aotearoa - New Zealand Qualifications Authority
- Te Pātaka Whaioranga - Pharmaceutical Management Agency
- Mahi Haumaruru Aotearoa - Worksafe New Zealand

Signatory agencies not interviewed

Some signatory agencies have minimal or no use of algorithms, so it wasn't appropriate to engage in a full interview. These agencies were sent a questionnaire.

- Te Manatū Mō Te Taiao – The Ministry for the Environment
- Te Tūāpapa Kura Kāinga – The Ministry of Housing and Urban Development
- Toitū Te Whenua – Land Information New Zealand
- Te Manatū mō ngā Iwi o te Moana-nui-ā-kiwa - The Ministry for Pacific Peoples
- Te Manatū Waka – The Ministry of Transport
- Te Kāhui Whakamana Rua Tekau mā Iwa—Pike River Recovery Agency
- Te Minitatanga mō ngā Wāhine — The Ministry for Women
- Te Ope Kātua o Aotearoa - New Zealand Defence Force

- Te Arawhiti — The Office for Māori Crown Relations
- Te Tari Arotake Matauranga — The Education Review Office
- Manatū Aorere — The Ministry of Foreign Affairs and Trade
- The Mental Health and Wellbeing Commission

B.4 Document review

As well as interviews and questionnaires, the review encompassed relevant documentation and literature. The review of documents was designed to supplement the interviews and questionnaires with background information to help answer the set of questions in the terms of reference. In particular, the review used relevant documents to:

- Find further case studies of impacts on agencies' data and algorithm governance
- Establish timeframes for how agencies have implemented the charter
- Understand broader perspectives on the safe and effective use of algorithms and how this could aid transparency of government agencies' algorithm use in New Zealand. Documents used are listed in Section 3.5.

'F
TAYLOR
FRY