



Centre for
Public Impact

A BCG FOUNDATION

Briefing Bulletin:

Analysing AI

Its risks, our recommendations

Artificial intelligence (AI) is an unmissable opportunity for government, but it has many risks. What are they and how can we respond to them?

Broadly, we think there are two risks for government when it comes to AI:

- 1- It could “do nothing”; or
- 2- It could “do wrong”

Each of these scenarios presents unique risks.

The “do-nothing” risk

• Why do we have governments?

We believe that government can be a powerful force for improving the lives of citizens. One of the key ways it does this is through the delivery of quality services (for example, justice, education and health systems).

Citizens yield certain liberties (through mechanisms such as taxation and a willingness to submit to the rule of law) in exchange for a government that delivers quality outcomes to all, regardless of their economic circumstances. This social contract rests on an assumption of legitimacy – governments are regarded as legitimate because they represent the interests of the populus rather than private individuals.

While some academics posit that the manner in which a government is formed (through elections, for example) is a powerful contributor to legitimacy, modern literature focuses on the need for government to deliver quality outputs. According to this theory, the greater the quality of outcomes, the greater the legitimacy of a government. Accordingly, if outcomes decline, so, too, will government authority.

• Declining outcomes, declining legitimacy

There are two reasons we believe that the quality of government outcomes may be at risk.

First, if governments have less **access to information** than other entities, their capacity to make quality decisions diminishes. There is a direct correlation between the amount of information one has to hand and the quality of the decisions one can make. In a world where the government knows less about a citizen than a private entity like Facebook does, we foresee issues with making quality decisions

Second if government lacks **the capacity to use knowledge and data** it will be outperformed by private entities that attract and retain talent. Governments have been collecting vast amounts of data on their citizens long before the private sector made the concept of big data popular. But most governments aren’t actually making use of these vast data stores. Not only is **public sector data collection haphazard but many**

government agencies lack the necessary in-house talent and infrastructure to maintain and analyse all the information they possess.

A lack of technical expertise in organising and analysing their data has also meant that **governments have had to pay for assistance from the private sector**. Not only are these missed opportunities for governments but there are potentially serious risks if governments persist in their lethargic uptake of AI.

We are already seeing government completely pulling out of offering services in certain areas and outsourcing service provision to private companies in others. For example, in 2016, a suburb of Tampa in Florida experimented with replacing two bus lines and subsidising Uber rides for its citizens instead. At US\$40,000 a year, the programme will be about a quarter of the cost of the two bus lines it is replacing.

Yet there are, and should be, serious concerns about what this kind of private sector service provision might mean. Three concerns that we have identified are: the regulatory and labour issues surrounding private sector companies’ service provision; the ethical concerns surrounding private sector use of government data about its citizens; and the technological deskilling of government staff, as data management is increasingly outsourced.

Yet, even if governments manage to react successfully to the regulatory, ethical and human capacity demands discussed above, an even bigger risk remains: **in a world where the private sector is better able to serve the needs of the public, what role remains for governments, and where will they find their legitimacy?** The balance in society between state and corporate power has long been complex, but a scenario in which public service provision is increasingly taken over by a more capable private sector will mark a substantial shift.

The do-wrong risk

The alternative risk we imagine is one in which governments do deploy AI but do so in ways that could be perceived as detrimental to the social good.

There are three potential issues for government in this situation. It may:

1. **Systematise inequity;**
2. **Encounter spiralling cost;**
3. **Abuse its power.**

1. Systematised inequity

Whilst in theory AI is noise-free and neutral, in practice it is often neither. **Data can be noisy** because of spurious readings, measurement error or

background data, resulting in randomly inconsistent and inaccurate outputs. Even unsupervised MLAs are dependent on the subjective decisions made by the humans who select and tag the data used to train them.

An example of this is the filtering of objectionable content in social media. To train an algorithm to identify such content, a human has to flag a series of cases that they find sensitive for reasons that may be based on personal or cultural norms. There is a risk that the veneer of objectivity that algorithms provide mask the very real subjectivity that lies underneath all data. Governments must be careful to ensure that marginalised communities are not further marginalised by algorithms that bake in assumptions about their circumstances.

2. Spiralling costs

Whilst cost reduction is one of the central reasons cited in framing the decision to introduce AI, a major risk to government is the hidden costs that AI adoption brings. **The assumption that AI delivers cost-savings is not universally true;** the costs of implementation are often indirect, difficult to calculate, and therefore underestimated. Take the US job-matching search engine for veterans. The search engine costs US\$5 million per year but is only used by a couple of hundred veterans. Government approaches to AI appear to assume a range of cost-savings, leading them to discount the need to measure and plan for costs, particularly lifecycle costs such as data and model management, oversight mechanisms, and auditing expenses.

Furthermore, **the risk is that once AI systems are implemented, it will be difficult to remove them. This is because they become embedded in complex processes** and there is a reduced human capacity to fulfil the role that people once performed. If this is the case, governments will be locked into particular systems, with the potential for spiralling costs.

3. Abuse of government power

A final and more ominous ethical concern is the potential for **government misuse of citizens’ data**. For governments to take full advantage of the benefits of AI, it will be necessary to share and centralise data across all of government.

Yet the risks associated with this are huge. On the one hand, the requirement for systems to be robust (for example, against hackers) will be even more acute than it is at present. On the other, there is also the need to be wary that government itself does not overextend its powers, turning into an **all-seeing “Big Brother”** state, monitoring and manipulating the minutiae of its citizens’ lives.

Even where governments do not move into new functional areas of control, the increased automation

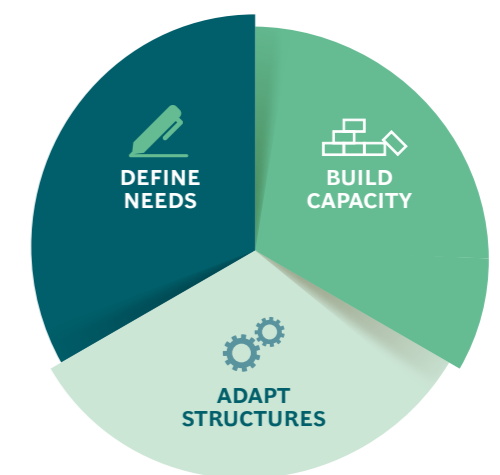
of decision-making through the use of AI may be perceived by the public as dangerous. Fears of faceless decision-making may be well-founded if algorithmic complexity is such that decision-making processes become inaccessible and incomprehensible to the average citizen. This means that transparency and accountability will be diminished. Moreover, these changes will most probably affect the least advantaged groups in society, who are both more likely to be subject to the outcomes of these decisions and less likely to understand the processes involved in producing them.

Ultimately, mitigating the ethical risks involved in using AI **requires people – and not technology – to be in control**. If AI is to be successfully implemented by governments, expertise and transparency is required at all levels. Governments need to ensure that they have the personnel to interact meaningfully with AI, while citizens need to be able to scrutinise the results of AI-based decision-making.

What should we do about it?

In light of these risks and challenges, we make the following recommendations to all government departments and officials:

- **Define needs:** establish best practices for identifying departmental needs.
- **Build capacity:** build the human and technical capacity necessary to the uptake of AI.
- **Adapt structures:** change the existing cultural, regulatory and legislative environments.



Further reading:

Want to dig deeper?

You can find more details on these recommendations in our AI Report:

- “AI in government” (brief 1)
- “The risks of AI” (brief 3)



Centre for
Public Impact
A BCG FOUNDATION

The Centre for Public Impact is a global not-for-profit foundation, funded by The Boston Consulting Group, dedicated to improving the positive impact of governments.

.....

Contact:

Adrian Brown

adrian@centreforpublicimpact.org

centreforpublicimpact.org

Join the conversation

Follow us [@CPI_foundation](https://twitter.com/CPI_foundation)  [in](#)  [f](#)

Learn more

Public impact is an evolving story.

Stay up to date via the Centre for Public Impact's Observatory of case studies from around the world at centreforpublicimpact.org/observatory

And learn how governments can improve the results they achieve for citizens by checking out the Public Impact Fundamentals at centreforpublicimpact.org/fundamentals