Centre for
Public Impact
A BCG FOUNDATION

TAXATION

ARTIFICIAL INTELLIGENCE
in government

Artificial intelligence in taxation

A case study on the use of AI in government

Summary

- This is a case study that describes a group of public sector AI interventions, in this instance AI applications in the field of taxation. It draws on several OECD countries' experiences of using AI for collecting and processing taxes, anonymised due to the sensitivity of the topic.
- These countries' tax authorities face a number of challenges: encouraging taxpayers to comply with the law, detecting and preventing illegality such as evasion and fraud, and reducing the requirement for seasonal workers in tax offices.
- They have focused on understanding AI's potential by modelling applications like insolvency risk assessment, whose cost benefits appear very favourable.
- This case study examines other promising areas for AI-based advanced analytics, for example identifying the cases most likely to lead to the recovery of revenue, and predictive communications with taxpayers.
- The authorities have demonstrated these AI prototypes' technical feasibility, but problems remain, such as concerns that by opening the code to scrutiny they may enable auditors and others to game the revenue system.

To learn more about how to proceed along a path towards using AI in the public sector effectively and legitimately, read our other AI case studies and our accompanying paper: How to make AI work in government and for people.

Advanced analytics

Taxes support all public sector functions. While the shape and nature of tax policy is highly political, it is less contentious to suggest that governments should collect and enable tax collection in the easiest and most efficient ways they can.

A government that does not waste money when collecting taxes, but instead employs effective mechanisms to find the individuals and organisations that are non-compliant and to raise money from them is in the general public interest. Being able to raise taxes effectively and efficiently enables government investment in education, infrastructure and welfare.

Many tax agencies are exploring the use of advanced analytics and AI to meet their two main public functions in the 21st century. The first is to manage tax compliance in order to detect and prevent illegal behaviour; the second is to provide services and education to taxpayers to help them comply more easily.[1]

Leading tax agencies have drawn upon analytic methods for over 25 years, particularly to meet part of the first function: to choose the right cases to audit.[2] It is impossible for any tax agency to control, manage and check every single taxpayer. It is also a waste of scarce enforcement resources to routinely examine low-risk, compliant taxpayers, not to mention an inconvenience that is likely to be unpopular at the ballot box. The use of analytics to manage these challenges has further increased in recent years, with 53 percent of OECD tax agencies reporting (in 2015) that they used predictive modelling in their detection and enforcement work.[3]

The scope of advanced analytics is now expanding into new areas, going beyond identifying suitable cases for audit to include methods of enforcing filing and payment compliance, providing better taxpayer service and debt management, and informing policy at a more strategic level.[4]

Many tax agencies share a generic model for identifying risk, analysing and prioritising it, taking action, and evaluating the outcome. Different types of analytics tackle different aspects of this model. **Predictive analytics** attempt to anticipate *likely problems*, so that tax agencies can take appropriate preemptive action. **Prescriptive analytics** help tax agencies model the impact of *actions* on taxpayers, so they can choose the most effective course of action for any given sector, segment or case.[4] New and promising approaches have recently linked different parts of this model together. In particular, tax authorities in Australia, Canada, Norway, the UK and elsewhere have started to build AI-supported models that attempt to predict which high-risk individuals are most likely to react positively to different tax authority interventions for recovering revenue.[4] These techniques, which draw on behavioural sciences as well as data mining, build on existing methods by ensuring that the cases representing the best value for the agency are pursued as a priority.

What makes AI a success in these contexts? This case study draws on interviews in different OECD countries, where interviewees talked frankly about their experiences building advanced analytics for better taxation. They have been anonymised due to national sensitivities.

AI initiatives in taxation

AI in the taxation domain consists of a wide variety of narrowly and sharply defined tools for particular tasks. In this section, we outline some of the applications of AI that may not be well known or expected.

Tax agencies can use AI systems to communicate differently with different groups of citizens for maximum impact. One analyst noted that if an individual "had a track record of not paying early, but their financial data said they were doing pretty well", then they should be picked up by a specially designed model that says "X has a capacity to pay but not a propensity to pay". In that case, rather than "starting with a weak letter, then moving a bit more firm and so on, [they] might jump to stage three of the intervention process and go harder earlier". Even low-risk scores can inform action. In these cases, "if you have a low risk score, then this might just encourage us to send an SMS, as we think you're going to pay, but we could increase efficiency on this group [of individuals]". This avoids spamming customers who would not respond to such approaches, and allows agencies to develop specific interventions tailored to a particular group. AI has opened up opportunities for new, simple policy options which were not possible before.

Tax agencies can also use AI to save money by not having to take on seasonal labour and temporary workers. Agencies often hire a large amount of seasonal labour to clean or fix millions of forms which are not correctly processed through the systems. One tax agency explained that while these workers tend to "look for errors and clear those items", this is possibly due to the increased use of both logic-based systems and machine learning by agencies in automating this process. They are "increasingly working out how to automate the fixing of these forms", and have even realised that the seasonal workforce "made a lot of errors, and [agencies] have managed to detect and correct many of these with new automation processes". The money tax authorities save in this way can then be invested in other tasks, such as making relations with citizens smoother, easier, and more human-centric.

As indicated above, AI also allows agencies to predict potential non-compliance to maximise revenue and fight tax avoidance. Models are often trained with particular sectors and groups in mind, and used across the organisation by auditors in their day-to-day roles. One agency recently invested in a large external project to build an insolvency prediction system for firms, illustrating the importance of better intelligence about business health rather than just about illegal practices. Some agencies are also experimenting with clustering methods, using AI systems to understand which taxpayers are most similar to each other, and to segment them in new ways. Social network analysis is helping to uncover complex VAT frauds that span multiple actors, and AI can be used to augment this as datasets become larger and more complex.[4]

Enabling practices and infrastructures

Individual tax agencies have their own enabling practices for designing and deploying advanced analytics and AI-powered tools. Flexible computing infrastructures are important in enabling agencies to build machine learning systems. One agency "developed [their] own Linux servers with root access, which we can use to deploy open-source software how and when we want". Another was trialling the powerful open-source statistical software, R. Yet a third emphasised that there were liability challenges when doing this, and that they "don't use R, mainly for reasons of commercial contracts" – meaning that there was no large company either to sue or to provide support.[1] [1]Although they had been talking to the Health Ministry, which does use R, they "internally see [their] models [as] more pivotal and operationally critical, [so] need to be able to phone up a supplier if something stops working". The flipside of this problem is that most graduates come out of university knowing free and open-source tools such as R, and converting that knowledge to proprietary tools is often costly, time-intensive and unattractive to them as an employment prospect. They see proprietary tools as much more limited, analytically weak, and a poor investment for their future careers.

Flexible computing infrastructure and in-house expertise allows tax agencies to engage high-tech vendors on equal terms. As one respondent noted, "a lot of vendors are 95 percent open-source software with a veneer or user interface on the top that's their own", which means that in practice the agency can use its in-house developers to "prototype a similar product on our [servers] and then talk to [large technology vendors], small startups or whatever, and better frame what we need and want". Having a proof-of-concept prototype in-house allows for greater efficiency, because vendors are aware that they are dealing with an organisation of rival capacity. However, working with vendors after that point is also of benefit, because the in-house developers can move on to creatively diagnosing

1 Such problems have been partially solved by large firms, such as Microsoft, developing commercial versions of open source software such as Microsoft R Open

and analysing further problems that might benefit from AI approaches.

Machine learning and other AI systems have often been built or overseen by specially designed teams that can review each other's work, sometimes with new specialist expert functions.

> *"We're usually working with a lead data scientist or modeller, and someone else who is a bit less involved, providing peer review along with colleagues in the rest of the team,"*

said one respondent. A notable innovation in one tax agency was the introduction of "team leaders with expertise in particular areas, roving reporters, so-to-say, in areas such as data science and predictive analytics". These team leaders do not have specific management responsibilities or teams, but instead can flexibly "engage with model-builders at different stages in the process, such as when output data is produced, or during implementation", providing much-needed mentorship and guidance, and transmitting knowledge across and within the institution, as well as going to different government departments to explain what the tax agency is doing.

Tax agencies have had to introduce agile processes in new and unique ways. One agency used "tiger teams [a concept in agile delivery] so we can actually get a model to piloting in perhaps a week". They cannot import this organisational style of working straight from the private sector, however, and had to design a way to integrate lawyers upstream in the development process in the same agile way, so ensure that if the agency "changes a tax form or alters a deduction, for example, [they're] legally allowed to do that".

Different agencies take different approaches to centralising or decentralising their analytics functions. Australia, Canada, New Zealand and the US have a decentralised approach, whereby analytics teams work closely alongside operational teams. Ireland, Mexico and Singapore all centralise their analytics teams in their head office. In the Netherlands, the analytics unit is separate, and can innovate quickly without being hindered by the slower rate of other departments. Other countries, like the UK, have a mixed setup.[4] While there is no clear "best" model here, managing that tradeoff between being operationally embedded close to "where the action is" and being free to innovate and experiment is a constant challenge, and tax agencies have to be sensitive about it.

Lessons and challenges

Tax agencies have often focused their efforts on understanding the potential gains from deploying an AI system. As many of these interventions are new and untested, they are usually uncertain in their benefits, and so the gains from developing a model (which might not work) will be unclear. One agency discussed how they create benefit cases for the models that are being deployed, in order to estimate the potential gains or savings. To make such a case, they used "historic information to inform hit rates, money recovered, and such like", because "model-making costs a lot of money, perhaps three to six months of time, and often new software or IT costs are incurred". In particular, they need to create models to understand how their new models are likely to save them money, such as "whether we will get yield back from these models within the next three to five years". In many cases, the potential benefits seem substantial: "in the case of expensive models such as the insolvency risk model currently being procured [by this agency], it is expecting around 10 times cost back in benefits," the agency noted.

Tax agencies have adopted different views on building explainable AI systems. One agency observed that they

had "better buy-in" when they provided the logic of the system to the customer, as it demonstrates "a lot of value-add, particularly where an administrative decision needs explaining to a customer, or it goes to a tribunal, or so on", and "helps reputation-boosting at the point of client relations". Another agency noted that they do not release many details, and "although they might say [they] are interested in sectors or size, and perhaps [might] share the weights with one or two key people", they are "primarily concerned that if the model weights were public, their usefulness might diminish".

A third case raised an interesting challenge. The tax agency was unwilling to share model weights or detailed explanations with auditors in case they gamed the system in question. This could happen in two main ways, according to the agency. Firstly, the auditors might use explanations of risk scores to find cases they could develop into a larger fraud, at which point they could let it accumulate and catch it afterwards. Under proposals by which auditors get a commission or in part earn their own wage, this creates a system of perverse incentives. Secondly, and perhaps more importantly, there were concerns that auditors would start making judgements based on the explanation rather than the risk level. If risk scores looked at particular cities, for example, and marked them as high risk, auditors might start searching for a disproportionate number of cases in certain regions. This would be problematic, because the model might be flagging a more nuanced correlation, so higher-level management might be unable to control issues of skew, bias and discrimination within the AI systems they deploy and manage.

Concluding remarks

Tax agencies have been at the forefront of deploying analytics in government, and AI is no exception. Agencies have benefited from interventions, such as automating the tasks for which they formerly used seasonal labour, to pursuing cases where they would be more likely to recover tax revenue, to making it easier for taxpayers to be compliant through predictive communications projects. This has necessitated new techniques and practices within the organisations themselves.

While there is no one-size-fits-all model, a few general lessons are evident. Analysts and designers are often located near frontline, operational staff, so they can discuss and tailor interventions appropriately. In-house capacity gives agencies the ability to mock up the analytics systems they want *before* they engage contractors, so that the problem definition clearly reflects the situation on the ground, and technical feasibility has already been demonstrated. New software tools are rapidly becoming available, with 11 OECD organisations beginning to use or experiment with powerful, open-source tools such as R, enabling them to hire and retain talent more effectively than with the less flexible software traditionally used in the sector.[4]

Challenges remain, though. Transparency is a double-edged sword – making a system too open may promote gaming by individuals seeking to evade the tax system. It may also reveal uncomfortable truths about areas or sectors that are neglected or disproportionately targeted by auditors.[5] Furthermore, analysing uncertain and often expensive AI design initiatives to understand whether they will provide sufficient benefit relative to required investment is also an area in need of further methodological development. A complex AI system has high risks but potentially high rewards – it might have huge benefits, or might not work at all. Yet these issues are not insurmountable, and new balances, methodologies and tradeoffs will continue to be created and navigated. Looking ahead, AI systems are set to continue to make tax easier to pay for citizens and fairer in its collection and enforcement for all.

BIBLIOGRAPHY

[1] Munawer Sultan Khwaja, Rajul Awasthi and Jan Loeprick, Risk-Based Tax Audits: approaches and country experiences, The World Bank, 2011

[2] OECD, Advanced Analytics for Better Tax Administration: Putting Data to Work, OECD Publishing, 2016

[3] OECD, Tax Administration 2017 - Comparative Information on OECD and Other Advanced and Emerging Economies,OECD Publishing, 2017

[4] OECD, Advanced Analytics for Better Tax Administration: Putting Data to Work (n2), OECD Publishing, 2016

[5] Michael Veale, Max Van Kleek and Reuben Binns, Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. ACM Conference on Human Factors in Computing Systems, 2018

Centre for
Public Impact
A BCG FOUNDATION

Join the conversation

ai@centreforpublicimpact.org
Follow us @CPI_foundation: #FindingLegitimacy
#AILegitimacy

October 2018
© Centre for Public Impact 2018