Chapter 2

Advanced analytics activities

This chapter describes how administrations are focussing their advanced analytics efforts across the range of operational activities, including audit case selection, filing and payment compliance, debt management, taxpayer service, policy evaluation and taxpayer segmentation. It also provides a high-level description and discussion of the main analytic approaches and techniques used in each area. It does not attempt to identify best practices (since these will vary widely depending on context), but it does highlight a number of significant applications and opportunities that administrations may wish to consider.

Advanced analytics activities by country

Table 2.1 provides an overview of where survey respondents that are actively using advanced analytics have allocated their efforts. The table makes no assessment of the relevant capability of administrations working in these areas; it seeks only to identify where work is being carried out.

Table 2.1. **Summary of activities by country**

	Audit case selection	Filing & payment compliance	Taxpayer Service	Debt Management	Policy
Australia	■	■	■	■	■
Canada	■	■	■	■	
China		■			■
Finland	■	■		■	■
France	■				
Ireland	■	■	■	■	■
Malaysia	■			■	
Mexico	■	■			■
Netherlands	■			■	
New Zealand	■	■	■	■	■
Norway	■	■	■	■	
Singapore	■	■	■	■	■
Sweden	■			■	■
Switzerland	■				
United Kingdom	■	■	■	■	■
United States	■	■	■	■	■

Source: FTA Advanced Analytics survey, 2015.

Advanced analytics for audit case selection

Survey responses indicate, perhaps unsurprisingly, that audit case selection is the principal application of advanced analytics techniques. Of the 16 administrations that responded to the survey and are actively utilising advanced analytics, 15 indicated that they had deployed analytics to prioritise cases for investigation, audit, or other compliance intervention.

The survey highlighted a number of issues and practices for consideration in case selection:

a) Value of large datasets

Combatting VAT fraud or error, both in payments and re-payments, emerged as a particular area of focus in case selection, with France, Ireland, Mexico, the Netherlands, Norway, Sweden, Switzerland, and the United Kingdom all making reference to analytical models built to identify claims which might be fraudulent or otherwise non-compliant.

Many administrations choose to begin their analytics efforts in VAT non-compliance, attracted by the plentiful data arising from the high volume of payment and re-payment claims. Data-rich environments such as this favour analytics for two reasons:

- Firstly, because "noise" in the data (movements that are due to chance rather than any underlying cause of interest) tends to cancel itself out in larger samples, leaving a clear signal which can be learned by a model;
- Secondly, because the volume of returns is such that direct examination by experts is typically not feasible.

b) Importance of assessing against next-best alternative

The latter raises an important point about predictive models: a key part of the decision to build a model is an assessment of the next best alternative. Where this alternative is under-developed, advanced analytics can add significant value without difficulty; where this alternative is mature, and particularly where it uses information that cannot be made available to a model (for instance, in the form of "local knowledge"), it can be very challenging to add value through predictive analytics. This should be borne in mind when selecting advanced analytics projects.

c) The role of social network analysis

In addition to building statistical models to predict VAT fraud or error, several administrations (including Ireland, Malaysia the Netherlands, New Zealand and Singapore) are carrying out social network analysis (SNA) to help detect VAT carousel fraud (a VAT carousel is a complex form of missing-trader fraud which exploits the VAT-free treatment of cross-jurisdictional sales) and other group-level risks.

SNA helps administrations to identify risky groups in situations where individual-level assessments may fail to detect anything of concern. It identifies links between individuals (for instance, through company directorships, joint bank accounts, or shared telephone numbers), and assembles connected individuals into easily visualised networks. Caseworkers can then browse these networks to profile individual risks. Equally, the networks can be scored

for risk using either a rules-based assessment or a statistical model trained on historical data.

Perhaps because SNA is still a relatively new approach, our survey found that rules-based assessment is most common. As the practice matures and historical data accumulates, it is likely that more administrations will build analytic models to determine network risk.

d) Power of fine-grained models

It is significant that at least two administrations deploy Social Network Analysis alongside predictive models for the assessment of VAT risks. This illustrates an important limitation of individual analytical models: it can be very difficult for one model to predict a range of different risks with accuracy.

This is because the underlying relationships in the data are likely to be different depending on the nature of the risk in question: for example, the patterns that predict an error on a VAT return are likely to be very different to the patterns that predict a VAT carousel fraud. When a single model is asked to recognise both patterns, its performance tends to suffer. For this reason it may be advisable – resources permitting – to build a different model for each distinct risk type. This is acknowledged in Norway, where the model that predicts fraud and error in VAT declarations is not used to detect carousel fraud.

Clearly this raises a question as to how best to define each risk type: for analytic purposes, the key question is whether the underlying patterns and relationships are likely to be substantially different for the various risk types under consideration. It is likely that this will be primarily a question for expert judgment.

A variety of administrations, including Ireland, Norway and Sweden, have developed models to assess income or payroll tax risk, including identifying the likelihood of fraud or error on relevant deductions. Interestingly, the Netherlands has developed a suite of models to address the issue. One model predicts the likelihood that a case will yield; a series of models beneath this predict the likelihood that individual elements of the return are fraudulent or otherwise incorrect. Another model predicts the likelihood that the case will yield over a pre-specified amount.

Although more time-consuming and resource-intensive to implement, this method can offer several advantages over the usual single-model approach: Firstly, it helps caseworkers to direct their efforts more effectively at the riskiest aspects of the case, thereby facilitating more efficient case-working and encouraging adoption. Secondly, if the underlying patterns that link input variables and risk differ across risk types, this approach will generate better predictive performance.

e) Role of unsupervised learning methods

The applications described above are generally examples of *supervised models*: these are models that seek to learn from historical data where the outcome of interest (e.g. whether or not a case was non-compliant) is known. Clearly these models must be built using data collected from previous audits or other interventions; the exercise then is to "train" a model to recognise the patterns that best predict non-compliance.

The upshot is that models of this type will typically refine and automate our existing understanding of risk. In general, the main contribution of a supervised model will be to reduce the number of cases wrongly flagged for intervention. This will save caseworkers' time and lessen the burden on compliant taxpayers, but in general will not help to identify new or previously unknown types of risk.

To achieve this, administrations typically need to create *unsupervised models* – that is, models that seek to identify interesting or anomalous patterns in the data, rather than trying to learn from the outcomes of specific cases. Survey responses provide two good examples of such models: Australia's *nearest neighbours* model, which is designed to identify incorrect income-tax deductions, and Ireland's *income-consumption* model, which aims to identify under-declaration of income. Although the two models use different statistical techniques (*k* in the case of Australia's *nearest neighbours* and multiple regression for Ireland's *income consumption*), both operate on the same intuition: by comparing a taxpayer's return to those of his or her peers, it is possible to identify outliers for further investigation, and also to identify cases which, though they may appear unusual on initial inspection, are in fact normal when compared to other, similar cases.

Which approach is more effective depends on the specific problem at hand: where the administration's main aim is to reduce nil-yielding interventions (eliminate false positives), *supervised models* are generally most appropriate; where the aim is to identify previously undetected types of non-compliance (eliminate false negatives), *unsupervised models* are typically more suitable.

f) Other case-selection projects of interest:

Other initiatives undertaken include:

- *Structured Income flows:* A model that links analysis of related entities to uncover misreporting at the entity-level and non-compliance associated with the structure of income flows (United States).
- *Tax Agent risk:* A predictive model to assess risk at the level of the tax agent, rather than just the individual taxpayer (Australia).

- *Unreported Income:* A predictive model to specifically identify unreported income, as distinct from over-claiming of deductions (Sweden).

Advanced analytics for filing & payment compliance

The objective in filing and payment compliance initiatives is either to secure an outstanding payment or return, or preferably to prevent the problem from occurring in the first place: in either case, the operational aim is to change taxpayer behaviour. To achieve this outcome administrations are applying both *prescriptive* and *predictive* techniques. Survey respondents indicate that *predictive* techniques are used to identify taxpayers who are likely to fail to meet their obligations, while *prescriptive* techniques are used to determine how to communicate most effectively with this group.

Survey responses showed that advanced analytics techniques are being successfully applied to improve both the timeliness and extent of filing and payment compliance, including:

Use of experimental designs

Experimental design is a prescriptive-analytics technique in which treatment and control groups are partitioned and observed in order to isolate the effects of specific actions, interventions, or treatments. Direct taxpayer communications are particularly well suited to this type of work, since the number of cases tends to be large and the cost of creating variations and partitioning treatment and control groups is minimal. Many survey respondents mentioned this type of work: the Norwegian administration, for example, has engaged with a behavioural economics researcher to test a variety of communications intended to improve compliance on declarations of foreign income.

Blending predictive modelling and experimentation

Tax authorities including Australia, Canada, Norway and the United Kingdom have implemented programmes of risk modelling and controlled experimentation that identify which cases are likely to fail to meet payment or filing obligations, and which interventions are likely to remedy the problem. In initiatives such as these, analytic outputs are used both to prioritise cases and to determine treatment paths.

The United Kingdom has commenced building models that assess taxpayer risk prior to filing. The models predict which taxpayers are most likely to miss filing deadlines, in order to target interventions to encourage compliance. These interventions are based on insights gathered from the

United Kingdom's Behavioural Insights Team, which was established to apply *nudge theory* (a mix of behavioural economics and social psychology) to try to improve government policy and services. A typical outcome of such upstream models may be to support a decision to communicate intensively (e.g. by phone) with a small set of taxpayers believed to present high risk, as opposed to using blanket communications which may be expensive and inefficient.

In general, filing and payment compliance initiatives offer interesting examples of the use of analytics to support business processes from end-to-end. Rather than using statistical techniques only to make predictions, or only to identify effective interventions, it is possible to combine the two tasks to design and evaluate intervention programmes that might otherwise be thought to be prohibitively expensive.

Box 2.1. **Canada's use of data mining models for non-filer programmes**

The Canadian tax system is based on voluntary compliance. In Canada more than 25 million individuals pay and file their tax returns without intervention. The Canada Revenue Agency (CRA) manages its programmes using a risk-based approach, to direct resources to cases with the highest risk of failing to file on time.

The CRA has developed and continues to refine several predictive models to assist in the delivery of its non-filer programmes. The models support improved workload selection and prioritisation for the programmes, and also supply estimates for cases that have not filed returns. In its first year in production, one non-filer model resulted in a total of CAD127.6 million in additional positive assessments. The CRA is now moving away from a pure predicted value to a relative ranking indicator, dynamically scoring accounts on an ongoing basis. The CRA has also developed several other models to improve programme effectiveness and enhance taxpayer services by predicting self-resolution and responsiveness to a specific compliance action.

In addition to predictive techniques, CRA applies prescriptive analytics to support improved strategic and operational programme delivery. Prescriptive analytics is used to enrich the CRA's understanding of the non-filer population, optimise operational processes, and direct the application of compliance activities, allowing for more fact-based decisions. Complementing the use of predictive models, the non-filer programme is expanding its use of behavioural economics through nudge experiments to influence taxpayer compliance behaviours.

Source: Canada Revenue Agency (CRA).

Advanced analytics for debt management

Survey responses indicated that most administration's approach debt-management analytics in the same way that they approach analytics for filing and payment compliance: applying a mix of predictive modelling and experimental techniques to identify which cases should be subject to intervention, and which specific interventions should be carried out.

A number of administrations commenced their work in debt-management analytics by modelling the risk that an individual or company will fail to pay: Finland, Ireland, Singapore and Sweden have all built models that attempt to assess the likelihood of insolvency or other payment problems. In a similar vein, Australia and Norway have built real-time debt management systems that put in place different payment arrangements depending on a taxpayer's predicted propensity and capacity to pay. The Australian Tax Office (ATO) also uses predictive analytics to send SMS messages to individuals found to be a payment risk. In relation to interventions, the Canada Revenue Agency has run controlled experiments to determine the impact of automated work processes and different taxpayer communications.

Traditional predictive modelling in debt management helps tax administrations identify groups of potentially high-risk debtors. While this approach helps focus resources on targeting those cases of highest risk, a number of administrations consider it possible to take the application of advanced analytics one step further by using a combination of experimentation and modelling to identify which cases are most likely to respond to a debt-management intervention. These may or may not be the highest-risk cases as identified by traditional predictive modelling – models may highlight risky cases that are not amenable to intervention or cases where debts would be re-paid even in the absence of intervention.

To help administrations identify the course of action that will yield maximum incremental return, a technique known as *uplift modelling* is required. This approach starts by running a controlled experiment to determine the incremental impact of a particular intervention. It then applies predictive modelling techniques to identify which types of taxpayer show the greatest response. This model can then be used as a basis for targeting future interventions. This approach has become increasingly common in the private sector in recent years. In the field of tax administration, it is most likely to be used in debt management, taxpayer service and programmes encouraging voluntary compliance. Survey responses indicated that, to date, minimal use has been made of these techniques in tax administration outside of a few isolated projects. Administrations should seek to explore this area over the coming years.

Advanced analytics for taxpayer service

To date, tax administrations have tended to use advanced analytics mainly to inform case selection, compliance activities, and debt management. However, as discussed in the OECD publication *Increasing Taxpayers' use of self-service channels* (OECD, 2014), many administrations have commenced using analytics in support of taxpayer service. The use of analytics to assist in developing views on taxpayer channel use, inform design decisions and to identify opportunities to offer self-services are assisting tax administrations improve outcomes. The use of pro-active messaging, calling, and other interventions, especially in the case of possible non-compliance has encouraged administrations to look more closely at how advanced analytics can more broadly improve service delivery to taxpayers. Survey responses highlighted a number of initiatives in this area, with more planned for 2016.

One particularly innovative service analytics project was highlighted by Singapore, where the text of incoming customer emails is mined in order to classify, analyse, and gain insight into the content of taxpayer inquiries. This insight is then use to devise and prioritise initiatives to improve service delivery. This initiative is the subject of a more detailed case study below.

New Zealand applies similar tools for sentiment analysis and question extraction, and has also worked to achieve a customer-centric view of its data, integrating customer complaints, survey results, and risk management data to offer a more rounded picture of each taxpayer. This provides a platform for new operational practices – for example, fully customer-centric data makes it possible to treat complaints from compliant taxpayers differently to complaints from their non-compliant counterparts.

A range of administrations – including Canada, Ireland, Norway and the United Kingdom – are using a mix of predictive and prescriptive analytics techniques to manage which channels taxpayers use for inbound communications. In general, analytics is used to encourage greater adoption of digital channels. This in turn is expected to open up new opportunities for analytics, since tracking and experimentation are simpler and less costly in the digital environment.

Box 2.2. **Text mining of inbound emails in Singapore**

In 2014, the Inland Revenue Authority of Singapore (IRAS) began using text-mining techniques to analyse the content of emails received from taxpayers. The findings from this analysis complemented existing analyses of structured data and helped the IRAS to pre-empt or reduce contacts and improve service delivery for taxpayers.

Box 2.2. **Text mining of inbound emails in Singapore** *(continued)*

The objectives of the project were to identify the nature of taxpayer inquiries and highlight important changes and trends that might require response. Text data from taxpayer correspondence was extracted, cleansed, and structured to derive patterns and insights. The project was an iterative process that required close collaboration between the analysts and the business users to contextualise the findings and improve the text mining process.

As a result of the project, the IRAS was able to uncover insights, otherwise locked in textual data, on issues pertinent to taxpayers. In one project, text-mining helped to identify the common queries taxpayers had after an existing tax policy was changed. Based on this analysis, the IRAS was able to push out appropriate campaigns in a timely manner, to provide more guidance on the IRAS website, and to proactively initiate updates to taxpayers, thereby reducing the need for taxpayers to contact the IRAS.

Ongoing tracking of the nature of email enquiries over time has also enabled the IRAS to identify trends in certain topics and respond accordingly. Text mining has now replaced the manual tracking of email enquiries, which has saved time and improved staff productivity. It has also enabled the IRAS to track the nature of enquiries more objectively, avoiding the inconsistencies of interpretation typical of manual tracking.

Source: Inland Revenue Authority of Singapore (IRAS).

Advanced analytics for policy evaluation

Although most analytics work is carried out to support operational decision-making, survey responses highlighted that analytics is also being used for decision-making in relation to strategy and policy. The most common analytic applications in this field are tax gap measurement, and assessing or forecasting the impact of changes in tax policy.

China, Finland, the United Kingdom and the United States all use analytic techniques to carry out tax gap analysis. Here, the focus tends to be on using yield data from random audit programmes, since these typically give an accurate representation of the wider taxpayer base.

Singapore has deployed visual analytics and simulation methods to explore the likely impact of proposed policy changes. The use of data visualisation has also enabled policymakers to quickly identify patterns, trends, and anomalies, improving the efficiency of policy review and decision-making. China's analytics function has carried out assessments of the impact of major

tax reform initiatives based on simulation modelling, as discussed in the case study below.

It is worth noting that a simulation model is quite different to the type of predictive model discussed above. Where predictive models fit patterns to historical data, simulation models tend to draw more heavily on modeller input allied to economic theory (specifically, mathematical representations of the macro-economy). These models can help policymakers to understand and explore complex relationships, but their predictions or other findings will only be as accurate as the inputs supplied by the modeller.

Box 2.3. **China's assessment of the impact of value-added tax reform**

In 2012, the Chinese government implemented a value-added tax reform pilot programme, which replaced the previous business tax (BT) with a value-added tax in selected sectors. In order to analyse the overall effect of the policy reform, including effects on the wider economy, on tax revenue, on industry structure, on social welfare, and on a wide variety of economic indicators, China's analytics department built a Taxation Computable General Equilibrium (CGE) model.

The Taxation CGE Model consists of six parts: the production sector, private sector, government sector, tax sector, trade sector, and macro close condition. The social accounting matrix table was established on the basis of the Chinese input-output table, national economic statistics, and tax statistics.

Under the changes made in 2012, the manufacturing, wholesale, and retail industries were made subject to VAT while other industries remained subject to the existing business tax. Where the business tax was replaced by VAT, the output price would vary because VAT and BT applied different administration and calculation principles. This price change is the source of the effects modelled. Through the CGE's price system, the model makes it possible to estimate the different effects of VAT and BT on output price. The model even mimics the invoice deduction method of VAT.

Through the Taxation CGE model, the analytics team was able to estimate the economic and social consequences of the VAT reform. Their report was approved by the Premier of the State Council, and played a key role in the policy reform process.

Source: State Administration of Taxation (SAT), People's Republic of China.

Advanced analytics for taxpayer segmentation

Many survey respondents reported an interest in using analytical techniques to segment their taxpayer base more effectively, though few dedicated projects have to date been carried out in this area by tax administrations.

A number of administrations (including Ireland and the Netherlands) have experimented with unsupervised segmentation techniques. These techniques, which fall under the broad heading of "cluster analysis", seek to identify groups of taxpayers who are similar to each other in some significant respects, and dissimilar to the other groups identified. These projects have often provided interesting general insight into the taxpayer population, but have typically not shown a strong practical impact as the segments identified have not had obvious business applications.

An alternative approach to segmentation, currently being pursued in Ireland, looks to group taxpayers based largely on their predicted response-to-intervention. If all taxpayers respond in the same way to a given intervention, then there is little practical value in segmentation; where there are large and consistent differences in response-to-intervention, then segmentation is worthwhile, and should follow the observed differences in response. This approach, which uses the *uplift modelling* techniques described above, is likely to create multiple segmentations – ultimately, each type of intervention may require a different segmentation of the taxpayer base. With tax administrations increasingly looking to personalise service and develop appropriate and timely interventions, further work in this area will be of prime importance.