



## OECD global forum on VAT

Day 2: Plenary Session 2

Topic: Artificial intelligence for compliance risk management in domestic and cross-border trade

### Overview

Before I start presenting the slides, I want to give you the context of our operating environment and why we decided to investigate the AI path.

- Within my area, 9 staff manage the x-border GST risk, which includes, risk management, data acquisition and compliance.
- The GST registration threshold is \$75K. The threshold is quite low and we potentially have a lot of small non-resident businesses that we need to consider.
- The data we get from the financial sector is large relative to our capacity. We cannot assess the relevance of all the data.
- We are confident that the large non-resident businesses and platforms are in the system. This is thanks largely to our relationship with BIAC and its members.
- Politically, the Australian small business community have very effective lobby groups. One of their key focuses is ensure that there is a level playing field between domestic and international businesses in the digital age.
- When I set the challenge for my team to use AI, it was not to set the world on fire, my goal was to get them to be AI literate; to build a capability platform to explore what AI can do. What I did not expect was for the team to grow their capabilities exponentially.

### (Slide 1) Opening

1. Thank you for the opportunity to share Australia's experience in using Artificial Intelligence for compliance risk management, a topic of growing global interest. My name is Adrian, and a part of my role is the risk owner of the GST risk associated with offshore merchants supplying LVIG and DPS into Australia.
2. Historically, we relied on manual analysis of large third-party datasets to identify offshore merchants with potentially unmet Australian GST obligations. This process was slow, error-prone, and resource-intensive.
3. In September 2025, we introduced AI-powered tools to our analytics processes as part of an AI exploration project purposed to reduce manual activities, improve output quality and quantity, and identify opportunities for broader application of AI across the ATO.
4. In today's presentation, I'd like to share with you the work we have done and benefits we have realised to date, which centre around efficiency gain and improved insights.

## (Slide 2) Our approach to AI augmentation

5. Overall, the analytics cycle for the international GST compliance program includes 5 key activities: data acquisition, data validation, data preparation, data analysis and strategic insights. This chart summarises the incremental improvements we have achieved to date and have planned for the next 18 months.
6. By June 2026, we would like to be in a position where a tried and tested AI framework can be applied beyond the international GST compliance program and across the ATO.
7. On this slide, we marked the manual activities in grey, staff-driven semi-automated activities in orange, and AI-driven fully automated activity in green. You will notice is that, not all the improvements are driven by AI, at least not now. This is because:
  - To fully realise the potential of AI, we need to explore ways to refine and automate existing processes by having the right policy, methodology and/or technology in place, and
  - Human interaction and judgment are essential to our role as the tax regulator.
8. I will now walk you through how each of the data analytics activity is improved in each phase of our AI journey to date.

### Data acquisition and validation

9. The data acquisition and validation stages set the foundation of our analytics cycle. They remain manual due to the need for engagement with external stakeholders – major Australian banks. Our engagement effort with the banks include:
  - Explaining the data requirements including the key information (e.g., merchant name, merchant ID, industries, countries and sales)
  - Explaining how we will use the key information (e.g., identify legal / controlling entities, contact them to advise GST obligations, and potentially raise assessments)
  - Providing a draft of the information gathering notice for the banks to review, including allowing time for the banks to test their own systems and confirm whether they can meet all requirements
  - Issuing the formal notice and receiving the raw datasets
10. Whilst the initial data gathering was time consuming, by building a positive relationship with the stakeholders and providing a clear and mutually agreeable set of requirements, subsequent data acquisition and validation time in 2025 was **reduced by 50% from 16 to 7 weeks**.
11. In addition to the above stakeholder engagement that requires manual intervention, the ATO has strict rules on how data of this nature – referred to as Special Purpose Acquisition Data (SPAD) – must be handled. Whilst the corporate policy could change in the future, right now data validation by an ATO officer is required to minimise errors and ensure downstream outputs are reliable and for the intended purpose.

### Data preparation

12. The data preparation stage focuses on consolidating and cleansing data from multiple sources – in our case four datasets from four different banks - to create a unified entity view of the merchants.
13. This process poses a significant challenge due to the different data standard and format by each bank. Often, the data includes multiple variations of the same merchant (e.g., PTY LTD vs. LTD), making it difficult to consolidate relevant merchants and sales under a single entity.
14. Historically, we relied on Excel to conduct data preparation, which involves significant manual examination and matching due to Excel's limited 'match' functionality. The manual nature also prevented us from interrogating large datasets or carrying out population analysis.

15. In 2024, we began trialling RStudio – a development environment for R, a programming language for statistical computing - to conduct data preparation. The methodology at the time relied only on merchant names and accounted for minor variations such as Pty Ltd. vs Ltd. After the consolidation using this methodology, the population size was **reduced by 51%**.
16. Just now in 2025, we refined the RStudio script by including more matching parameters such as emails, phone numbers, industries and countries. This improved methodology reduced the **population size further** by a further 9% to **60%**.
17. Overall, the transition from data preparation by Excel, to basic RStudio scripting in 2024 and the current, advanced RStudio scripting, has led to high-quality data for downstream analytics and **reduced the data preparation time by 87.5% from 8 weeks to 1**.
18. In 2026, we will continue to optimise the methodology by introducing more parameters to the RStudio script and testing the auto-run of pre-defined script through Copilot.

## Data analysis

19. The data analysis stage transforms the prepared dataset into actionable insights by identifying a pool of candidates for potential compliance treatment based on certain risk indicators and profiling each candidate to determine whether compliance treatment is warranted and if so, what is the risk hypothesis and appropriate treatment option.
20. Historically, the candidate pool was less sophisticated. We ranked all unique entities from the consolidated datasets by sales and manually profiled the largest 600 entities by conducting desktop research by extracting public information such as the email and postal address, legal name, key contact, phone number, and terms and conditions of each offshore merchant. This exercise took 10 weeks to complete.
21. In 2025, we began using Copilot for profiling, where the scanning and categorisation of public information and generation of profile documents were automated. Staff were then required to exercise judgment to validate the AI-generated outputs. Whilst maintaining the same FTEs and number of entities profiled, the AI-driven process **reduced the time commitment from 10 weeks to 3.5 hours**.
22. The time saving allowed us to develop more sophisticated candidate pools underpinned by more complex risk indicators such as sales increase between 2023 and 2025 and comparison between the sales and reported GST of a registered entity. I'll share an overview of our candidate pools in a moment.
23. We have observed that Copilot delivers more accurate and complete profiling outcomes when working with up to 20 entities per prompt. Beyond this threshold, the results became inconsistent between entities and often fell short of manually generated equivalents. This is likely caused by the limit in Copilot's processing power compared to other AI solutions such as the enterprise versions of ChatGPT or Google Gemini.
24. Going forward, we will explore opportunities to enable bulk profiling, potentially integrated with AI-powered tools such as Fivecast and NexusXplore.
  - Fivecast is an open-source intelligence platform (OSINT) that collects and analyses social media, websites, chat forums, and even the **dark web** and generates deeper entity and trend analysis.
  - NexusXplore is designed to **link people and entities**, which is useful in identifying entities that change business names to avoid detection.

## Strategic insights

25. This final stage generates insights that inform key population trends. Historically, the analysis was done at a sub-population level (e.g., the top 600 entities) largely due to the time required to manually conduct the data preparation and analysis activities.

26. The realised benefits to date, including the accuracy and efficiency gain from data preparation and profiling<sup>1</sup>, has allowed us to develop a methodology to segment the population into multiple risk categories, which will support tailored treatment options.
27. Staff are also able to conduct deeper population analysis to uncover strategic insights such as merchant churn, new entrants and exited businesses, sales trends, and shifts in country of incorporation. These insights open new opportunities for compliance intervention.
28. Looking ahead, we will integrate Copilot into above population analysis for automation and test its capability to generate risk treatment strategies, determine risk ratings, and potentially select cases. Our objective, whether it is with the current AI solution or not, is to move into generative AI.
29. Lastly, the section at the very bottom 'Beyond the horizon' captures a broader enterprise-wide future state where we move beyond Copilot by building core AI infrastructure capabilities tailored for targeted ATO taxpayer, legal and risk-specific datasets. This vision is in its initial stage and aimed at scaling AI throughout the ATO to increase compliance, assurance, and confidence in the tax system.

### **(Slide 3) Strategic insights (1): population analysis**

30. Before wrapping up today's presentation, I just want to share with you some of the strategic insights that this new data analytics process has helped us achieve. Whilst some of the insights were still generated by staff members, we are in the process of testing the potential of AI in generating similar, if not better, strategic insights, in an automated manner.
31. Notably, AI has helped us identify
  - 2,800 unregistered entities with potential Australian GST obligations, totalling \$2.9 billion AUD sales across four key industries – gambling, gaming, adult entertainment and crypto exchange. These entities will be the focus of our treatment strategy.
  - 40 registered entities with potentially unmet Australian GST obligations, totalling over \$2 million underreported GST.
32. We further broke down the 2,800 unregistered entities into the following risk categories with tailored treatment options.

### **(Slide 4) Strategic insights (2): registration risk**

33. We developed four colour coded risk categories, with red requiring one-to-one treatment and green benefiting from one-to-many treatment. These categories are:
  - Red: New entities or those with rapid growth and significant sales (over \$5m in 2025FY)
  - Amber: New entities or with rapid growth and sales between \$3m to \$5m in 2025FY
  - Yellow: Entities with sales between \$1m to \$2m in 2025FY
  - Green: Entities with sales below \$1m in 2025FY and growth over \$500k 2023FY to 2025FY
34. Compliance treatment against these entities will fully commence in April 2026.

<sup>1</sup> The improved process (current) includes advanced scripting to group the population, using Copilot to identify legal names for all merchant groups, matching the legal names to our registrant list (to identify registered vs. unregistered population).

## (Slide 5) Reflections

35. In conclusion and upon reflecting our AI-journey over the last few months, there are four things I'd like to share with you.
36. AI is a seed of an idea. You start with a simple task, and soon you might see an explosion of ideas and use cases across the organisation.
37. Be direct and specific, both in terms of the objectives of your AI initiatives and prompts into your AI tools.
38. Be open-minded about what AI can and cannot do. Depending on the policy, culture and risk tolerance of your organisation, consider the degree for human judgment whilst capturing the opportunities around agentic AI and automate the end-to-end process.
39. Most importantly, not every efficiency enhancement involves AI. The AI implementation requires us to re-examine our ways of working by removing and/or digitalising existing manual activities, areas that we might have overlooked prior to the AI journey.
40. Lastly, be honest with the current state by considering the capability, process and policy gap that needs to be addressed.