



[\(https://ia.acs.org.au/\)](https://ia.acs.org.au/)

# Centrelink trials AI to detect fraud, prioritise debts

## Controversial move prompts calls for transparency.

By Jeremy Nadel on Feb 27 2025 09:37 AM

The agency responsible for administering Centrelink has revealed few details about the artificial intelligence models it trialled late last year to predict fraudulent welfare claims and which debts to prioritise recovering.

A Services Australia spokesperson told *Information Age* that systems were “being thoroughly tested” before “deployed into ongoing business processes”, but welfare rights advocates have called for the release of audits of the systems’ accuracy and risk of bias.

“Services Australia should be required to provide evidence that the system does what it is supposed to do,” IT professional Justin Warren, who [won a seven-year Freedom of Information \(FOI\) case with the agency for the release of documents related to the RoboDebt scheme](https://www.theguardian.com/australia-news/article/2024/jun/07/robodebt-court-documents-it-expert-justin-warren-ntwnfb) (<https://www.theguardian.com/australia-news/article/2024/jun/07/robodebt-court-documents-it-expert-justin-warren-ntwnfb>), said.

The use of AI was “trialled to enhance work allocation to staff, specifically debt-related activities that are likely to result in a ‘finalise no debt’ outcome”, according to [documents released to an anonymous FOI applicant in August](https://www.servicesaustralia.gov.au/freedom-information-disclosure-log?context=22).



Centrelink trialled an AI system to detect fraud and which debts to prioritise recovering. Photo: Shutterstock

Another “predictive model” was developed that “compliments existing fraud strategies by assisting fraud analysts in identifying high risk claims.”

## The human-in-the loop debate

The FOI documents said that the models were “non-decision making” and the spokesperson added they do not “make decisions on whether a claim is fraudulent or not” or “determine any debts or complete any debt actions”, which suggests a system more like Robodebt’s predecessor: Online Compliance Intervention.

Welfare recipients’ reported income was being compared to their averaged annual tax records to automatically identify discrepancies since 2001, but the discrepancies were manually assessed, until Robodebt launched in 2016

(<https://ia.acs.org.au/article/2025/corruption-watchdog-to-investigate-robodebt-referrals-.html>), which sent debt letters without pre-checks.

However, digital rights advocate Asher Wolf said the use of the term “non-decision making” could be highly misleading.

“Calling a predictive model ‘non-decision making’ is an oxymoron, it is a contradiction in terms.

“Predictions always involve decision-making: which theory to investigate, which target to pursue.”

## AI debt prioritisation

The Services Australia spokesperson said that the “debt prioritisation pilot...involved a process to make workload prioritisation more efficient.”

“By prioritising these potential debts, staff can review them and action them right away.”

“The pilot is designed to assist staff; customers do not interact with the system.”

Both tools are defined in the FOI document as “machine learning”, suggesting their rules were set by feeding historical data to self-learning algorithms.

This contrasts with the pre-coded rules of the tools used in the RoboDebt scheme that raised more than half a million inaccurate Centrelink debts (<https://www.legalaid.vic.gov.au/learning-from-the-failures-of-robodebt>).



(<https://www.pearson.com/en-au/professional-personal-development/campaigns/acs-australian-computer-society-member-discount.html>).

## AI fraud detection

Services Australia's Fraud and Corruption Control Plan 2025–2026

(<https://www.servicesaustralia.gov.au/sites/default/files/2025-01/14400b-2412.pdf>), said last month that its “detection processes...include: advanced data analytics including machine learning to detect potential fraud and corruption.”

However, the spokesperson said that the model has not been “deployed more broadly” since its “proof of concept trial” last year.

“The model flags claims with potential risks to be reviewed and analysed by one of our trained fraud analysts.”

“A fraud analyst will determine if a claim requires further investigation for potential fraud, otherwise the claim will be processed under usual business processes.”

The spokesperson declined questions about what the model associated with fraud, whose data it was trained on, and the data types it was trained on.

“To ensure the integrity of our operations, we don't provide specific information on how our detection capabilities work.”

Wolf said that more transparency was needed to ensure that Services Australia's systems would not “automate...past injustices.”

“To protect against human rights abuses, algorithmic decision-making in high-risk sectors – such as human services – should be transparent, open-sourced, open to public scrutiny and have easily identifiable and accessible avenues for remedy if errors occur.”

“Are Services Australia's machine learning models choosing to target geographic regions that have historically been subject to high levels of welfare surveillance of Indigenous communities?”

“Or areas subject to the disastrous ParentsNext program or punitive Cashless Welfare Cards? Or will the algorithms directly target individuals who have previously received erroneous Robodebts?”

Services Australia built a machine learning tool in 2014

(<https://classic.austlii.edu.au/au/journals/IntJlCrimJustSocDem/2016/9.html>), which has since been discontinued, that “mined ten years of customer data in the Centrelink Mainframe...to identify demographic characteristics associated with welfare claimants that had been overpaid.”

However, the spokesperson said the current tool is used to detect fraud whereas the previous one was applied to the broader categories of non-compliance and debt recovery.



## JEREMY NADEL

Jeremy Nadel is a Melbourne-based freelancer whose work has been featured in *The Guardian*, *The Saturday Paper* and *Crikey*. He was most recently a staff reporter for *ITnews* and *CRN* reporting on the tech channel and enterprise IT. In 2023, he was awarded Best New Journalist at the Australian IT Journalism Awards. For story tips, you can email Jeremy at [securetips.org@protonmail.com](mailto:securetips.org@protonmail.com) (<mailto:securetips.org@protonmail.com>).

### Information Age Comment Policy

We welcome respectful feedback. Harmful comments, URLs and spam will not be published.

Got it

### Tell the Information Age team your reaction to this story...

4 Responses



Upvote



Funny



Surprised



Angry



Sad



Enjoyed

1 Comment

Login ▼

G

Join the discussion...

LOG IN WITH

OR SIGN UP WITH DISQUS ?

Name



Share

Best Newest Oldest



goju1

3 hours ago

One thing I know about government IT Systems (having used them and been involved in their development) - they are all sub standard with regards to contemporary standards. They are clumsy, error prone and counter intuitive.

So, when I hear about government developing AI based systems, I cringe with fear...

As an example, in WA, the government decided to upgrade a Patient Administration System that

