

Submission to

Consultation on Guidance for Transparency in Automated Decision Making

on behalf of

Australian Association of National Advertisers

24 June 2026

Introduction

The Australian Association of National Advertisers (AANA) welcomes the opportunity to provide this submission to the OAIC's Consultation on Guidance for Transparency in Automated Decision Making. As the peak body representing Australia's advertising and marketing ecosystem, the AANA is fully committed to protecting children's digital safety and privacy online.

Summary

The introduction of Australian Privacy Principles (APP) 1.7–1.9 establishes a mandatory framework for disclosing the use of personal information in computer programs that make, or substantially assist in making, decisions that significantly affect individuals. While primarily aimed at high-stakes sectors like finance and insurance, the broad definitions of "computer program" and "significant effect" heavily intersect with the automated ad-tech pipeline and personalised advertising. If the guidance sets thresholds too low, it will capture routine digital operations that Parliament did not intend to regulate, producing disclosure obligations that overwhelm rather than inform consumers.

While transparency enhances consumer trust by increasing awareness, a rigid regulatory application to real-time ad tech could also inadvertently compromise anti-fraud mechanisms and disrupt the legitimate utility of data-driven, personalised marketing.

This submission aims to flag issues around negative impacts on anti-fraud mechanisms and legitimate personalised marketing.

Meaning of "computer program" and "significant effect"

Impact on Personalised and Targeted Advertising

The digital advertising ecosystem relies fundamentally on automation, including OpenRTB (Real-Time Bidding), predictive audience segmentation, and algorithmic delivery optimisations. Under the new rules, the disclosure obligation is enlivened if an ADM system utilises *personal information* to make a decision that can *reasonably be expected to significantly affect the rights or interests of an individual*.

Advertising Below Significant Effect Threshold: The Guidelines should make clear that standard, everyday contextual or behavioural targeting used to serve general consumer advertisements should not meet the "significant effect" threshold. In addition, where most programmatic advertising operates on pseudonymised data, or signals where the individual is not reasonably identifiable, we note that it would be exempt from APP 1.7.

More fundamentally, standard advertising delivery (including audience segmentation, frequency capping, and algorithmic optimisation of ad placement) does not constitute a "decision" within the meaning of APP 1.7. Consistently with how the term "decision" is interpreted in the Privacy Act, where it carries its ordinary meaning, informed by its use across administrative law, corporations law, and the Administrative Review Tribunal Act 2024, a "decision" is a discrete, operative determination directed at an individual that produces a definitive outcome (granting/refusing a benefit, imposing a condition, issuing or revoking a licence). Standard advertising delivery processes generate continuous, probabilistic rankings of content for each user. They do not produce discrete, operative determinations directed at an individual. No individual is denied access to goods or services by virtue of which advertisement they are shown. AANA believes that OAIC guidance should confirm that content personalisation and algorithmic recommendation, including in the advertising context, do not meet the "decision" threshold since they do not directly result in denial of access to essential goods or services.

Supply Chain Accountability: The Issues Paper indicates that entities should "actively identify, assess and keep oversight" of third-party ADM usage. Automated ad delivery relies on third-party vendors and platforms. AANA seeks clarification on the practical expectations of such monitoring, including the frequency, depth and evidentiary standard required to demonstrate compliance.

The Compliance Burden of Shared Control: Programmatic advertising relies on a distributed chain of automated decisions involving brands, demand-side platforms (DSPs), supply-side platforms (SSPs), and publishers. Pinpointing whether the platform or the advertiser holds the principal disclosure obligation under APP 1.8 creates substantial legal ambiguity. AANA seeks clarification on what constitutes reasonable and proportionate steps to meet 'arranged for' obligations under APP 1.7(a), particularly where decisions are driven by third-party platforms such as demand-side platforms (DSPs) and other ad-tech providers. Guidance should confirm that entities may rely on vendor disclosures, contractual assurances and platform governance frameworks, recognising that detailed visibility into proprietary algorithms is not available.

Privacy Policy Inflation: If the OAIC interprets "substantially and directly related" intermediate steps broadly, businesses may be forced to include exhaustive, hyper-technical descriptions of their ad-tech stacks. This risks overwhelming consumers with dense disclosures, running counter to the primary objective of APP 1 (open and transparent management). The Issues Paper emphasises balancing meaningful transparency with usability. A layered approach is a practical solution already recognised in AI and privacy guidance globally and could be an appropriate solution to manage complexity while maintaining understandable disclosures.

Significant Effect, Cumulative Impacts: Under the legislation, the 'significant affect rights or interests' threshold must be attached to a discrete 'decision'. Both the qualifier 'significantly' and term 'decision' must retain meaningful limiting force in accordance with Parliament's intention. AANA believes that the Guidance should confirm that the obligation attaches to discrete decisions, which have direct, material, impact on an individual's rights or interests and not to cumulative and iterative effects over time (such as repeated targeting, optimisation cycles or sustained engagement). Given that digital marketing ecosystems rely on continuous optimisation and multiple low-impact interactions, guidance should confirm that ordinary cumulative engagement does not, in itself, constitute a 'significant effect' absent material harm, exclusion or disadvantage.

Influence vs Determination 'Substantially And Directly Related': AANA also seeks clarification on the threshold in APP 1.7(a) relating to whether a computer program is 'substantially and directly related' to making a decision. Guidance should distinguish between systems that determine decisions and those that merely inform, optimise or support human decision-making processes. Many tools (e.g. analytics, segmentation and optimisation systems) provide inputs into decisions but are not determinative of outcomes and should not be considered 'substantially and directly related' unless they are a key factor driving the final decision.

Inferred Data: AANA further seeks confirmation that inferred or derived data, including behavioural profiles, audience segments and predictive scores, is not generally personal data and not subject to disclosure. These forms of data are central to modern advertising and may influence ad placement and effectiveness without directly identifying an individual.

The Online Safety Benefits of Targeted Advertising

While regulatory attention frequently focuses on the data collection aspects of targeting, digital advertising frameworks leverage these exact automated systems to act as a critical shield for public interest and consumer safety.

A central tenet of modern, responsible advertising self-regulation, such as the AANA Codes, is ensuring that marketing is contextually and audience-appropriate. Far from being purely extractive, automated targeting systems provide necessary mechanism to actively exclude children and vulnerable populations from encountering age-restricted or inappropriate commercial material.

Key Digital Safeguards Enabled By Target Filtering

Targeting is essential for a tailored and safe online environment for children and other vulnerable people.

Target/Vulnerable Group	Hazard Prevented	How Algorithmic Filtering Achieves It
Children (Under 16/18)	Exposure to age-restricted products (Alcohol, Wagering, Occasional Food & Beverage/HFSS products).	Employs robust age-gating, behavioural signals, and structural exclusions to ensure campaigns for adult-only categories are strictly ring-fenced away from minors.
Individuals Experiencing Vulnerability	Financial distress from predatory targeting or exposure to triggering content.	Uses demographic and interest suppression lists to ensure high-risk financial products or alcohol promotions are not served to audiences showing indicators of financial hardship or recovery.
General Online Audiences	Invasive or malicious content.	Automated ad-review algorithms continuously scan and filter creatives before deployment, protecting users from harmful or deceptive scams.

Preserving the Mechanics of Protection

If transparency guidelines become overly restrictive or inadvertently penalise the data-processing required to build effective exclusion profiles, the capacity of platforms to shield vulnerable users will be compromised.

For instance, to ensure a child under 16 does not see an ad for an age-restricted product, the automated ad delivery system must process age markers or reliable behavioural proxy data to say *"Do Not Serve."* If the data minimisation or transparency requirements create a chilling effect on maintaining these exclusion lists, the industry's ability to provide age-appropriate experiences would face severe technical friction.

AANA is seeking the following from the OAIC guidance:

1. **Clarity on "Significant Effect":** The guidance should confirm that content personalisation and algorithmic recommendation, including in the advertising context, do not meet the "decision" threshold since they do not directly result in denial of access to essential services.
2. **Proportionality over Text-Heavy Policies:** Disclosure rules that favour clear, plain-language consumer summaries rather than demanding highly technical breakdowns of

complex programmatic ad bidding stacks, which would ultimately degrade the user readability the APP 1 amendments seek to achieve.

Impact on Ad Fraud Mitigation

The digital advertising ecosystem has a massive underlying vulnerability - ad fraud. Globally, digital ad fraud is projected to steal over $100 billion in ad spend, meaning that roughly one in every five ad dollars is swallowed by bots, malicious publishers, or sophisticated data manipulation rather than reaching real humans.

While transparency in algorithmic decision-making is a valid public policy goal, ad-tech platforms must deploy highly dynamic, opaque, and instantaneous automated defenses to analyse, flag, and filter out these evolving, multi-billion-dollar fraud tactics. Forcing anti-fraud algorithms to disclose their internal verification rules or variables under the guise of transparency would hand a complete blueprint to bad actors, allowing them to optimize their bots to bypass defenses entirely, directly endangering the integrity of the digital economy.

Ad fraud (e.g., botnets, click farms, and domain spoofing) is combatting complex, automated adversarial behaviour. Ad-tech platforms deploy sophisticated machine learning and rules-based ADM systems to identify, flag, and block fraudulent traffic in real time.

As regulatory frameworks push for more automated decision-making (ADM) transparency, understanding the mechanics of ad fraud and the catastrophic business impacts that make its elimination non-negotiable is critical.

The Core Problems: How Modern Ad Fraud Operates

Ad fraud has evolved far beyond simple, clumsy bot clicks. Today, it is an industrialised, AI-driven criminal enterprise operating at three distinct levels of the advertising funnel:

1. **Impression-Level Fraud (Fake Views)**
   Fraudsters exploit Cost-Per-Mille (CPM) campaigns by manufacturing fake impressions on ads that a human eye never sees:
   - **Pixel Stuffing:** Compressing an entire advertisement down into an invisible 1×1 pixel frame. The ad technically renders and bills the advertiser, but it remains entirely hidden from the user.
   - **Ad Stacking:** Layering dozens of ads directly on top of each other in a single ad slot. Only the top ad is visible, but the fraudulent script fires a billable impression to the ad networks for every single ad in the hidden stack.
   - **Domain Spoofing:** Deceptive publishers modify ad tech bid requests to disguise their low-quality, automated, or illegal websites as premium, high-CPM news outlets (e.g.,

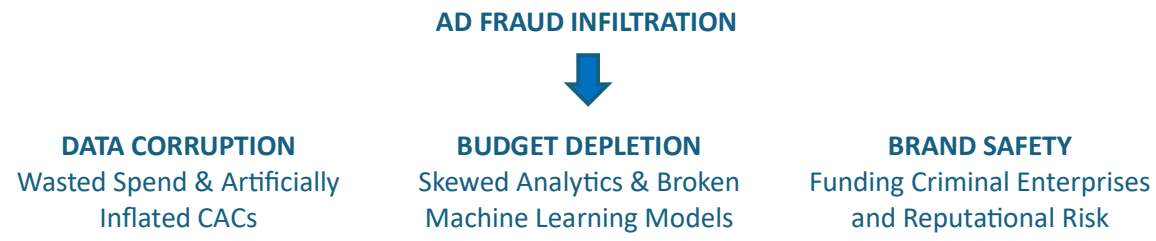
claiming the ad is running on a major national news site when it is actually rendering on a completely fabricated bot site).

2. **Click-Level Fraud (Fake Engagement)**
   Targeting Cost-Per-Click (CPC) budgets, this fraud rapidly drains ad spend while producing zero real business pipeline:
   - **AI-Powered Botnets:** Sophisticated bot networks utilize residential proxies to continuously rotate IP addresses. They use Generative AI to mimic human behaviours, such as cursor jitter, natural scrolling speeds, and varying dwell times, allowing them to slip past standard, legacy detection heuristics.
   - **Click Spamming & Farms:** Generating massive volumes of invalid clicks automatically (or via low-wage click farms) to exhaust a brand's daily advertising budget early in the morning, successfully boxing legitimate competitors out of the ad auctions.

3. **Attribution and Event Fraud (Stealing Credit)**
   This is the most financially damaging form of fraud because it targets the lower funnel (app installs, sign-ups, and purchases):
   - **Organic Hijacking / Click Injection:** Fraudulent software embedded in rogue apps monitors a device's activity. When it detects that a real user is about to organically download an app or complete a purchase, the fraudster quickly injects a fake ad click a millisecond before the event completes. The advertising attribution model mistakenly rewards the fraudster with a payout commission for a customer the brand would have acquired for free.
   - **Intent Mimicry:** AI fraud agents fill out lead-generation forms, create synthetic user accounts, and even add items to carts to make their traffic look exceptionally high-performing, triggering major performance payouts from brands.

Why Ad Fraud Must Be Eliminated

Eliminating ad fraud is not merely about recovering lost media spend; it is a fundamental requirement to protect the economic logic, data integrity, and safety of the digital business world.



- **It Corrupts Corporate Data and Analytics**
  Modern marketing relies heavily on data-driven feedback loops. When bot traffic floods analytics dashboards, it pollutes the data pool. A channel that looks like it has an incredible Click-Through Rate (CTR) might actually be entirely driven by bots. Marketers unknowingly double-down on toxic channels and cut funding to legitimate publishers with lower but human engagement.
  The ADM Risk: When ad fraud data trains downstream machine learning models and automated bidding systems, the AI learns to optimize campaigns toward behaviors that look human but are actually just highly optimized bot patterns.
- **It Destroys Return on Investment (ROI) and Ramps Up CAC**
  Because a significant portion of digital inventory is eaten up by non-human traffic, Customer Acquisition Costs (CAC) are artificially inflated. Legitimate consumers are blocked from seeing ads because daily limits are quickly swallowed up by malicious automation. Advertisers end up paying vastly more capital to acquire a single genuine customer.
- **It Subsidises Serious Criminal Enterprises**
  Ad fraud is highly lucrative, presents low entry barriers, and carries minimal risk of criminal prosecution compared to traditional cyberattacks. The billions of dollars siphoned out of corporate marketing budgets globally frequently fund transnational organized crime, ransomware syndicates, and digital bad actors.
- **It Threatens Brand Safety and Consumer Trust**
  A significant portion of ad fraud occurs via Made-for-Advertising (MFA) websites and platforms designed purely to harvest programmatic ad dollars using scraped or extreme content. When automated ad auctions fail to filter these platforms, a reputable brand's ads can appear next to disinformation, hate speech, or explicit content, instantly shattering consumer trust and brand equity.

The Unintended Vulnerability of Transparency

- **Exposing the Defensive Playbook:** Forcing platforms to explicitly detail the "kinds of personal information" and "types of decisions" automated anti-fraud systems make could provide bad actors with a blueprint. Malicious entities can reverse-engineer these disclosures to bypass detection thresholds, directly undermining ecosystem integrity.
- **The Mischaracterisation of Anti-Fraud Decisions:** Anti-fraud tools routinely make automated decisions to block traffic, filter IP addresses, or restrict ad delivery based on behaviour patterns. While these protect commercial interests, over-regulation could misinterpret these automated actions as "significantly affecting the interests" of individuals, even if the "individual" is an automated bot masquerading as human traffic.

Conclusion

To maintain a balanced approach that preserves both privacy transparency and economic functionality, AANA would support the following principles being incorporated into the OAIC's final regulatory guidance:

- **Establish a High Threshold for "Significant Effect" in Advertising:** The guidance should confirm that "significantly affect rights or interests" requires a discrete decision with genuine and material consequences for an individual's rights or interests (such as access to employment, financial products, healthcare, housing, or essential public services). Standard advertising operations, including behavioural targeting, frequency capping, and programmatic bidding, do not produce such outcomes. They do not deny access to goods or services, affect contractual rights, or produce legal effects.
- **Proportionality over Text-Heavy Policies:** Disclosure rules that favour clear, plain-language consumer summaries rather than demanding highly technical breakdowns of complex programmatic ad bidding stacks, which would ultimately degrade the user readability the APP 1 amendments seek to achieve and have negative implications for platforms' safety and integrity.
- **Proportionate Treatment For Security/Integrity Systems:** AANA recommends that guidance allow for high-level disclosures regarding the use and purpose of fraud, cybersecurity, security, safety and integrity systems, without requiring detailed disclosure of operational rules, thresholds or variables that could compromise effectiveness.
- **Commercial-in-Confidence Protections:** Reaffirm that disclosure requirements do not compel entities to reveal proprietary algorithms, machine learning models, or commercial-in-confidence data practices that would diminish competitive advantage or that are used for internal organisational purposes.
- **Clarity On Inferred Data:** Confirmation that inferred or derived data, including behavioural profiles, audience segments and predictive scores, is not generally personal data and not subject to disclosure.
- **Clarity on Chain of Responsibility:** Provide definitive safe harbors or clear frameworks outlining which party is responsible for fulfilling the APP 1.7 obligation.

Further Consultation

The AANA would welcome an opportunity to discuss in more detail with OAIC the issues raised in this submission. Please contact [redacted] at [redacted] regarding opportunities for further consultation.