



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Angus, D;Obeid, AK;Burgess, J;Parker, C;Andrejevic, M;Bagnara, J;Carah, N;Fordyce, R;Hayden, L;Lewis, K;O'Neill, C;Albarran-Torres, C;Cellard, L

Title:

ADM+S Working Paper Series: The Australian Ad Observatory Technical and Data Report

Date:

2024-01-08

Citation:

Angus, D., Obeid, A. K., Burgess, J., Parker, C., Andrejevic, M., Bagnara, J., Carah, N., Fordyce, R., Hayden, L., Lewis, K., O'Neill, C., Albarran-Torres, C. & Cellard, L. (2024). ADM +S Working Paper Series: The Australian Ad Observatory Technical and Data Report. ARC Centre of Excellence for Automated Decision Making and Society.

Persistent Link:

<https://hdl.handle.net/11343/340477>

License:

[CC BY](#)



ADM+S WORKING PAPER SERIES

> The Australian Ad Observatory Technical and Data Report


January 2024

Authors:

Prof Daniel Angus
Dr Abdul Karim Obeid
Prof Jean Burgess
Prof Christine Parker
Prof Mark Andrejevic
Julian Bagnara
A/Prof Nicholas Carah

Dr Robbie Fordyce
Lauren Hayden
Dr Kelly Lewis
Dr Christopher O'Neill
Dr César Albarrán-Torres
Dr Loup Cellard

WORKING PAPER 009



Acknowledgement of Country

In the spirit of reconciliation, we acknowledge the Traditional Custodians of country throughout Australia and their connections to land, sea and community. We pay our respect to their elders past and present and extend that respect to all Aboriginal and Torres Strait Islander peoples today.

Suggested citation

Angus, D et al. (2024) The Australian Ad Observatory Technical and Data Report, ADM+S Working Paper Series 2024 (9), ARC Centre of Excellence for Automated Decision-Making and Society.

Copyright © 2024 Daniel Angus, Abdul Karim Obeid, Jean Burgess, Christine Parker, Mark Andrejevic Julian Bagnara, Nicholas Carah, Robbie Fordyce, Lauren Hayden, Kelly Lewis, Christopher O'Neill, César Albarrán-Torres, Loup Cellard. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited.

ARC Acknowledgement

This research was conducted by the ARC Centre of Excellence for Automated Decision-Making and Society (CE200100005), and funded by the Australian Government through the Australian Research Council



Australian Government
Australian Research Council



EXECUTIVE SUMMARY

This report provides an overview of the Australian Ad Observatory project's data, infrastructure, and software tools, and serves as a detailed technical companion to the Background Paper previously published (Burgess et al., 2022). The Australian Ad Observatory ('Ad Observatory') is a project of the Australian Research Council (ARC) Centre of Excellence for Automated Decision-Making and Society (ADM+S) that aims to improve the observability of targeted online advertising in Australia using novel data donation infrastructure and hybrid digital methods.

The first section of the report provides an overview of the Ad Observatory project's data donation infrastructure, analytical dashboards and researcher tools. It describes, visualises, and discusses in detail the pattern of reported demographic characteristics and observations gathered from our pool of participants. We explain how this pool is positioned in relation to benchmark populations, with particular focus on the Australian Facebook user population. We outline the data types gathered from Facebook using this infrastructure, which include ad images and text as well as metadata such as the WAIST (Why Am I Seeing This) explanations provided to users. We also explain how meaningful information is extracted and aggregated from this data using the Observatory's tools and dashboards.

The report details the project's research questions and how these questions are pursued within nested case studies, each of which deploys purpose-specific but linked analytical strategies using the Ad Observatory's central infrastructure and tools. These questions are linked to methods that are uniquely enabled by the Ad Observatory, and that are designed to investigate the overall volume, dynamics, and targeting patterns of Facebook's ad system including the relationships between its temporal dynamics, demographic targeting, and the symbolic features of ad content.

The final section of the report details three case studies that have been undertaken within the project to date, including accounts of how they have worked with Ad Observatory data to answer topic-specific questions. These case studies are focused on areas of societal risk, where advertising is both potentially influential and subject to public oversight and regulation. The case studies include: investigations of environmental or 'green' claims in ads; the temporal patterns, sequencing, and targeting of alcohol advertising; and the extent and nature of online gambling advertising in Australia.


The report concludes with brief discussions of how the project team has approached key ethical and compliance challenges, including copyright, human research ethics, defamation, and the amplification of potentially harmful content.

KEYWORDS: PLATFORMS, SOCIAL MEDIA, ADVERTISING, DATA DONATION



TABLE OF CONTENTS

1.	The Australian Ad Observatory Project	6
2.	Overview of Project and Infrastructure	6
3.	Participants and Demographics.....	9
	Recruitment	10
	Challenges	10
	Counting Participants	11
	Power Law Distribution.....	11
4.	Ads	18
	Ad Observations.....	18
	Facebook Media Objects	18
	Calls to Action.....	19
	Advertisers	21
	Trends in Observation	22
5.	WAIST (Why Am I Seeing This Ad) Data	23
	Relationships Among WAIST Tags.....	24
6.	Research Questions and Case Studies.....	26
	6.1 Research Questions.....	26
	6.2 Case Studies	26
	6.3 Political Advertising.....	27
	6.4 Financial	28
	6.5 Scams	29



6.6	Green Claims	30
6.7	Alcohol Advertising.....	35
6.8	Gambling Advertising.....	37
7.	Ethical and Compliance Considerations	42
	Ethics and privacy.....	42
	Copyright.....	42
	Defamation and Allegations of Illegality or Misconduct	43
	Re-publishing Illegal or Misleading Content in Ads	43
8.	References.....	44



1. THE AUSTRALIAN AD OBSERVATORY PROJECT

The Australian Ad Observatory ('Ad Observatory') is a project of the Australian Research Council (ARC) Centre of Excellence for Automated Decision-Making and Society (ADM+S) that aims to improve the observability of targeted online advertising in Australia using novel data donation infrastructure and hybrid digital methods. To date, the project has focused on the Facebook advertising system, but the approach could be adapted for other platforms.

This Technical and Data report:

- provides an overview of the Australian Ad Observatory project's infrastructure and tools;
- describes its pool of participants;
- details its research questions, data, and analytics; and
- discusses some of the case studies that have been undertaken within the project, including accounts of how they have worked with Ad Observatory data to answer topic-specific questions.


The document is intended to work as a detailed technical companion to the Australian Ad Observatory Background Paper (Burgess et al., 2022), and can be read in combination with discussions of the project's methods or analyses drawing on the project's data. We anticipate it will also be useful to colleagues seeking to establish data donation infrastructure for the study of online platforms, as well as to the field of online advertising research. Unless otherwise stated, the data provided is for the 2022 calendar year.¹

2. OVERVIEW OF PROJECT AND INFRASTRUCTURE

Launched in October 2021, the Ad Observatory addresses two particular characteristics of online advertising that make it very difficult to monitor in the public interest: first, online advertising is ephemeral (ads appear and disappear and are not publicly archived); and second, it is personalised (ads can only be observed by the individuals to whom they are targeted).

The project overcomes these challenges by taking a citizen science approach and enlisting the help of research participants drawn from members of the Australian public who engage with the

¹ Actually 6 December 2021 to 6 December 2022 for technical reasons.



research project by contributing data donations (Araujo et al., 2022). Participants join the project by installation of a browser plugin that is available on a variety of leading web browsers (Figure 1). On Facebook, data donations take the form of ads observed on a participant's Facebook newsfeed. The plugin collects these ads along with their metadata and anonymous demographic information supplied by the relevant participant. Collected ads can then be reviewed from a personal dashboard offered through the plugin (Figure 2).

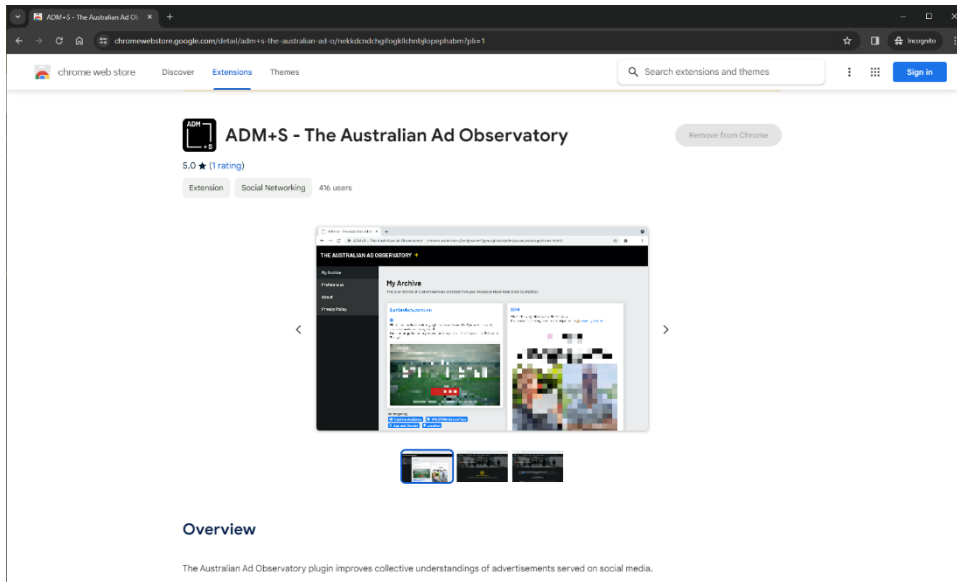


Figure 1. Plugin installation in a browser extension store

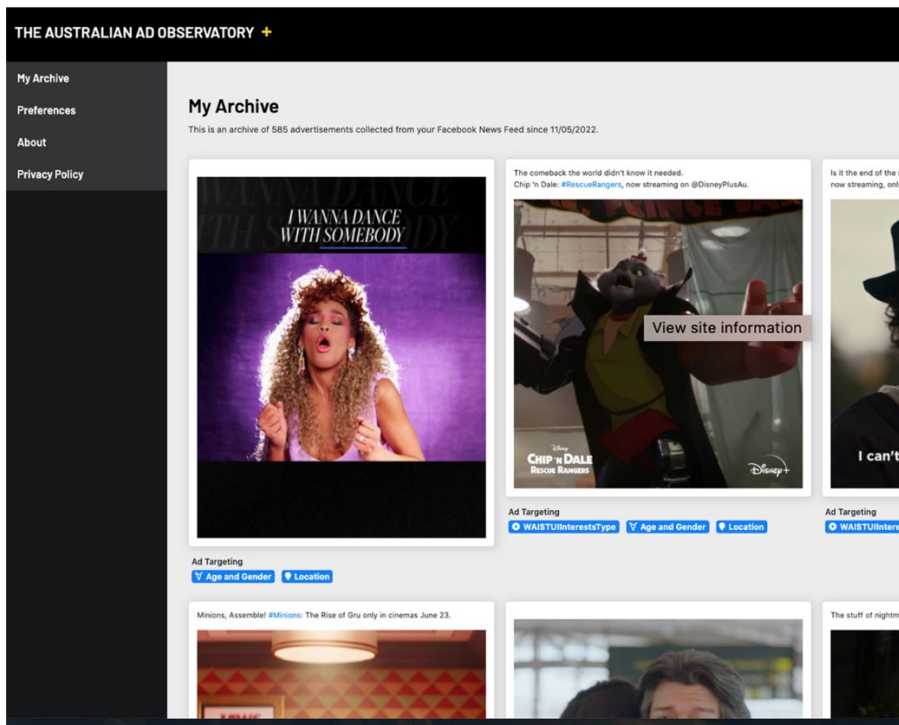


Figure 2. Personal dashboard of collected ads (as seen from a participant's perspective)



The project offers a separate research dashboard that brings together the various data elements of the ads collected by all participants to date (Figure 3). These data elements include:

- broad demographic categories of the participants to whom the ads were presented (and then 'observed' by the participant);
- time range of the observation;
- Facebook WAIST ("Why Am I Seeing This?") interest codes used in targeted advertising;
- ad metadata;
- text content.

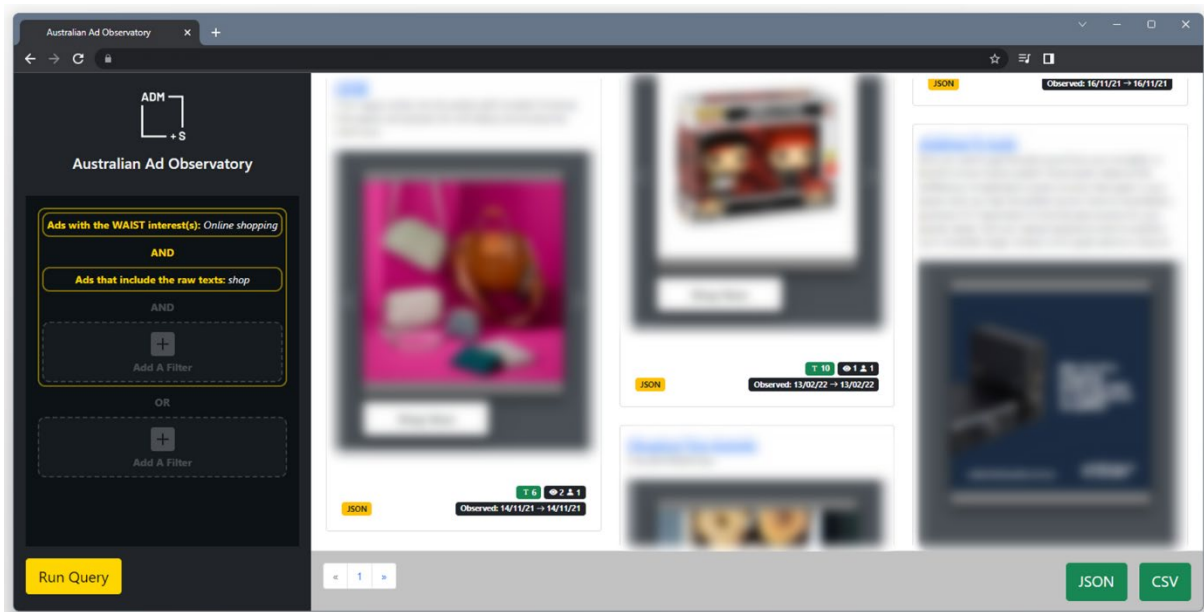
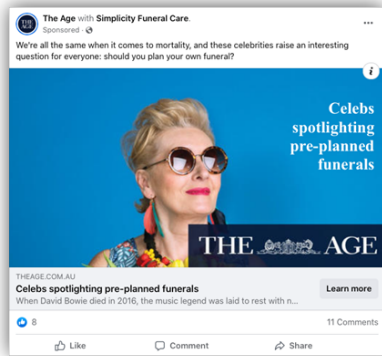



Figure 3. The second version of the research dashboard, accessible only by approved members of the research team

The static image content or video content contained within the ads is processed using automated techniques including Optical Character Recognition (OCR) text extraction or You Only Look Once (YOLO) logo detection (Figure 4). These techniques generate additional metadata that can be used to identify and cluster ads into categories or topics. Approved researchers can use the research dashboard to filter, search, and sort through these ads on the basis of different data elements. Filtered datasets can then be downloaded for offline analysis and interpretation.



"Celebs spotlighting pre-planned funerals"

OCR

THE AGE

Logo detection

Figure 4. Visualisation of Optical Character Recognition (OCR) and You Only Look Once (YOLO) image analysis techniques executed on the image content of an ad

3. PARTICIPANTS AND DEMOGRAPHICS

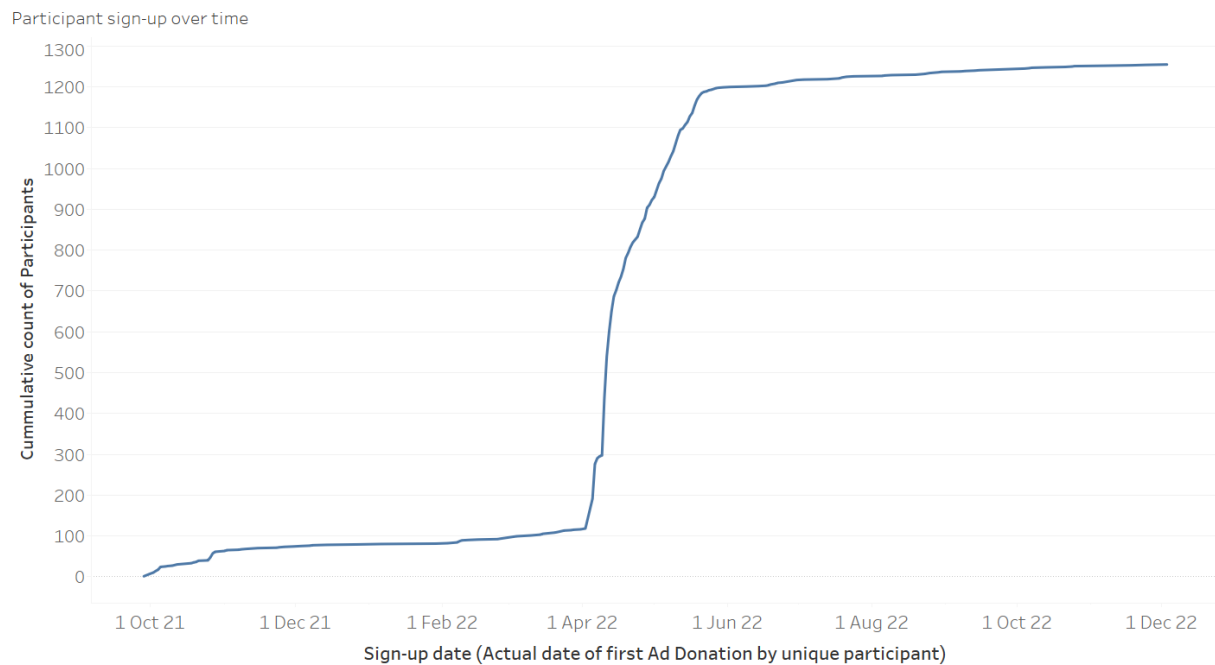


Figure 5. Participant sign-up over time for the first year of the project, indicating the date of the first observation made by a participant after successfully installing the plugin

Note: The total is around two thirds of the total sign-up of 2021 participants given that one third of participants never made a data donation.



Recruitment

Since launching in October 2021, the Ad Observatory has attracted a total of 2,012 participants, as of March 21, 2023.² There were 1,945 participants as of December 6, 2022.³ Recruitment between late 2021 and early 2022 was mostly via organic (i.e. unpaid) social media promotion, and organic promotion by the project team amongst professional and social networks. During April and May 2022, a series of nationally significant media appearances on the Australian Broadcasting Corporation (ABC) and Channel 7's breakfast television program *Sunrise*, plus a series of stories with direct links to the plugin during the ABC's Australian federal election coverage, meant the number of participants grew steeply within a matter of weeks. The timing of sign-ups to the project strongly coincided with the publication of a series of articles referring explicitly to the project on the ABC News website. We therefore make the assumption that a significant number of participants came from this ABC audience. After this growth point, we have sustained a slow trickle of new signups.

Challenges

The plug-in is designed to identify advertising within Facebook's site, which is something that Facebook has historically gone to great lengths to prevent. Facebook's approach has impacted important precursor projects by New York University and ProPublica, from which we drew part of our inspiration (NYU Cybersecurity for Democracy 2024, ProPublica, 2020). Over the life of the Ad Observatory project, Facebook has made a number of changes to their website. At times, these changes impaired the ability of the plugin to detect advertising on the website page.

Additionally, while roughly two thirds of our participants were able to make donations, one third seem to have been presented a different version of Facebook's page for unknown reasons. For participants who received this different version, our plugin was unable to detect any advertising. As a result, a third of our participants were not able to make any donations despite signing on to do so. This factor, and other factors (such as time spent using the online browser, uninstalling the plug-in, and using different devices to access Facebook) resulting in inconsistent donations, are important to take into account in downstream analysis. To avoid counting inactive participants, we presented sign-ups using the date of first donation in Figure

² Excluding Beta testing accounts. Unix Timestamp: 1679360400

³ The date for most of the statistics reported in this paper; see above note 1.



5. This provides an accurate snapshot of participant sign-ons only for those that were able to actively make donations.

Counting Participants

It is potentially misleading to only report on participant demographic quotas in terms of unique participants given that some participants may have donated very few ads, while others may have donated several hundreds or thousands of ads. We therefore cannot assume that donation rates are evenly distributed for all combinations of demographic qualities. We consequently avoid statistical analysis of participant demographics. Instead, we opt for reporting that takes into account the volume of donations per demographic. In discussing the data gathered from these donations in the following sections of the report, we refer to each instance of an ad being captured by our plugin as an 'observation', and in this context we refer to participants as 'observers'.

Power Law Distribution

The uneven nature of donations can be seen in the power law distribution of donations by unique participant, as presented in the histogram in Figure 6. This histogram reveals that a few participants donated a significant number of ads, while a long tail of participants donated progressively fewer ads. This distribution reaffirms the potential danger highlighted above in describing our participant base by unique participant characteristics alone, rather than taking account of the number of observations or donations. Analysis of the content of ads will need to take caution in making assumptions about how the prevalence of advertising themes and content relates to the demographic breakdown. It is, of course, possible to apply corrections and deploy statistical methodologies to build representative snapshots across select demographic groups in our subsequent work on the data.



Histogram of donation counts by unique participants




Figure 6. Number of ads donated by unique participants, organised as a histogram, note that the majority of participants (1,052) donated between 1 to 1000 ads, while a single participant donated 25,788 ads.

In providing a broad quantitative description of the data gathered to date, we chose the unit of ‘observation’, rather than ‘individual’ for the majority of our reporting. Each appearance of an ad is an ‘observation’. Specifically, we present the total count of observations made or donated by members of specific demographic categories, rather than a total count of unique participants alone.


As the main and most effective referral point for the project was the ABC’s coverage in April and May 2022, the audience demographics that we observe in our Ad Observatory participant population largely reflect the population characteristics of the ABC. The ABC Digital News population is typically middle to older-aged, is politically left-leaning, and has a tertiary education (Park et al., 2022). These same characteristics are seen in our population breakdown, as shown below. Additionally, we acknowledge that the mode of participation (a web browser plugin deployed on a desktop or laptop device) also favours an older demographic, as younger audiences are more likely to engage with news and social media via a mobile device (Park et al., 2022).

As we do not have access to Facebook’s internal data on user demographic characteristics, we rely on third party reporting by Statista (Hughes, 2023). Statista reports annually on gender and age characteristics of the users of various social platforms and media services (reproduced in Figure 7). Despite the characteristics noted above, a comparison of the Ad Observatory participant population characteristics and the age and gender breakdown provided by Statista



(Figure 8) highlights that we achieved a fair approximation of gender and age distribution for Facebook’s Australian audience. However, there was an under-representation of female users under 45 years of age and an over-representation of male users over 55 years (Figure 9).

As mentioned, the Ad Observatory participant population skews towards more highly educated with a significantly higher number of bachelors and post-graduate degree qualified participants (Figure 10) than what was observed in the 2021 Australian national census (Figure 11). For income distribution, we note a more diverse representation (Figure 12) in line with current national trends regarding income distribution. The political-leanings of the population suggests a more progressive-leaning population (Figure 13) as compared to the first-preference votes at the 2022 Federal election. This echoes a commonly heard refrain in citizen science projects, that they struggle to find participants who may hold more conservative political views (Bradley et al 2022). In terms of geographic distribution (Figure 14), we note a good representation across all states in Australia, largely reflecting population concentrations in major capital cities but with undersampling of rural and some regional areas.



© Statista 2023

Figure 7. Statista ground truth of breakdown of age and gender of Facebook users in June 2022 (the midpoint of our sample)

Note. From Hughes (2023).

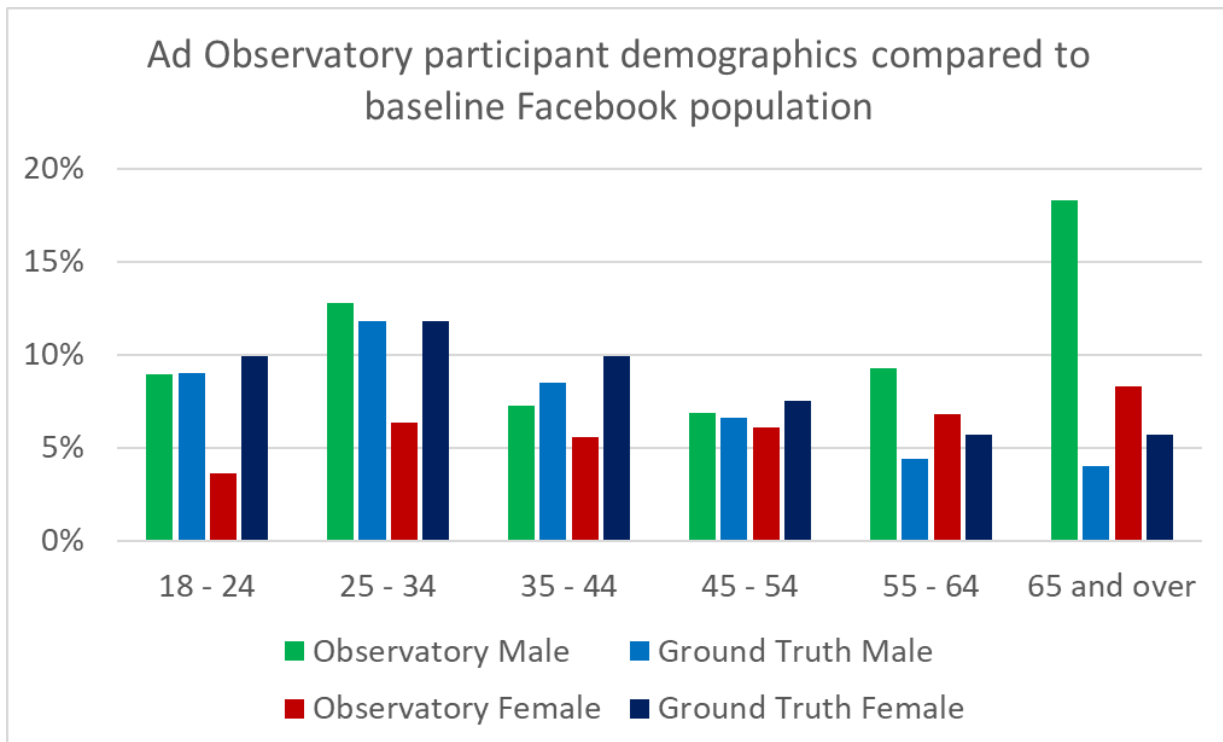


Figure 8. Combination of the Statista percentage breakdowns and Ad Observatory percentage breakdowns

Note: Adapted from Hughes (2023). This Figure omits the under 18 category of the Statista ground truth data as we did not have participants under the age of 18. The chart aggregates our 65-74 and 75+ categories to match the Statista ground truth categories. Our data largely follows the Statista ground truth distribution in the 65-74 and 75+ categories. Note that there is undersampling of females 18-24, 25-34, and 35-44, and oversampling of males 55-64 and 65+. For the purpose of this comparison only we also omit the 'other' and 'prefer not to say' gender categories from our data as these categories are not included by Statista. The 'other' and 'prefer not to say' gender categories account for 17,743 observations, and are included in all other analyses unless noted otherwise.




Figure 9. Age and gender distribution of Ad Observatory participants by total observation counts of demographic group

Note: This is not a tally of the number of participants, but a representation of the total number of observations by participants with these demographic characteristics.




Figure 10. Number of observations by educational group

Note: This is not a tally of the number of participants, but a representation of the total number of observations within these demographic groups.




Figure 11. Ground truth highest educational qualification of Australian adults as of 2021

Note: From ABS Census (2021). The educational profile of the Ad Observatory participants (see Figure 10) is biased towards those with bachelor and postgraduate degree qualifications when compared to the general Australian population (see Figure 11).




Figure 12. Employment versus income distribution by total observation counts of demographic group

Note: This distribution demonstrates that employment and income are loosely correlated as would be expected (colloquially the characteristics of the sample pass the 'sniff test').





Figure 13. Total observation counts of demographic groups by party political alignment according to 1st preference data

Note: From Australian Electoral Commission, (2022). The 2022 federal election results were: Labor = 33%, Greens = 12%, Liberal = 24% (Coalition = 39%), and Other = 17%.

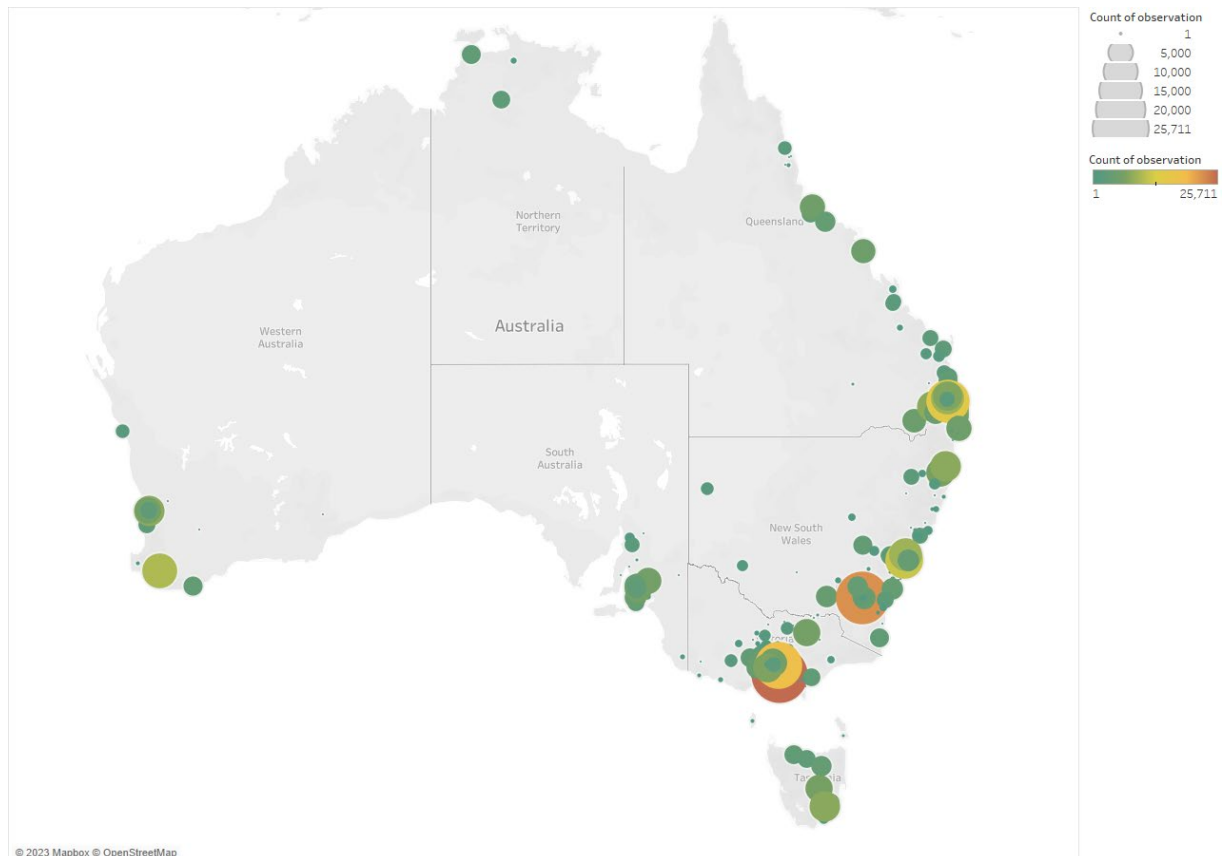


Figure 14. Geographical distribution of participants based on observations count

Note: The total number of observations are proportional to population at a state level. However there is an underrepresentation of participants in specific regional and rural areas (for instance, central and western QLD and central Australia).



4. ADS

Ad Observations

For the purposes of this report, we have limited our analysis to observations captured within the 2022 calendar year.⁴ In 2022, a total of 728,925 observations were collected. Of these, 279,450 were distinct advertisements. As demonstrated in Figure 15, the frequency of observations varies throughout the year, with the highest volume of observations occurring during the peak of the sign-up activity in May 2022, approaching the Australian Federal Election. The volume of observations also varies across each week of the year, with fewer observations occurring during weekends compared to weekdays. Significant drops in activity (around July 2022) and gaps in collection (between September to October 2022) are due to Facebook obfuscation efforts which blocked or hampered our ability to collect data from participants in these time periods. Obfuscation is a process that Facebook uses to prevent ad blockers within browsers from being able to detect ads within their webpage. As our plugin needs to identify ads within a page in order to donate them to the Ad Observatory, such changes to Facebook's webpage can result in periods of time where our development team has to update the plugin to once again identify advertising and enable donations.

Facebook Media Objects

Ads on Facebook predominately employ visual content, which may include image(s) , video(s), or a combination of both. Ads may use a single media object (either image or video), or include a carousel with multiple media objects as a set that users can interact with, usually by side scrolling. Within the observations, there were 47,374 carousels. Carousels comprised 6.5% of total, or roughly 1 in 15 ads encountered on the platform. Of the total observations during 2022, a simple count of the media objects used in the ads captured reveals that 30% of the observations included a video, either as part of a carousel or as a standalone media object (see Figure 16).

⁴ Observations were actually captured between 6 December 2021 to 6 December 2022 (due to a break in the browser plug in).




Figure 15. Frequency of ad observations (top) and active ad participants (bottom) in the Ad Observatory during 2022

Note: A participant is considered active if they have donated at least one ad during any given day




Figure 16. Number of observations that include a video, either as part of a carousel or as a standalone media object, in the Ad Observatory during 2022

Calls to Action

179,098 of the 728,925 ad observations included a discernable call to action (CTA) button, or set of CTA buttons for ads constructed as carousels. Many more ads likely contain a CTA. However, the ability of our ad donation infrastructure to identify all instances of CTA ad contents actually shown to participants is limited. This most likely relates to the obfuscation of ads within the browser that was discussed previously. While these CTA may appear as simple buttons to an end user, the situation behind the scenes (in the plugin code) is anything but straightforward. A significant part of our effort in running the Ad Observatory has not just been in identifying advertising, but also deconstructing the ad itself once it is donated.



In our efforts, we have recorded numerous instances of text contents within CTA buttons being meticulously jumbled, or rendered as SVG HTML elements, rather than appearing as conventional text elements. These, among various other formatting irregularities, have impeded attempts at systematically scraping the HTML content responsible for rendering CTA buttons. We identify this as an obfuscation technique similar to the kind previously described when attempting to identify ‘Sponsored’ content within Facebook newsfeeds. However, where the search criteria used to identify whole ads is close-ended (i.e., by identification of the ‘Sponsored’ term), the search criteria used to identify CTA buttons is open-ended and thus unpredictable, which has impeded the data reformatting process.

For those CTA buttons we have been able to capture and reformat, these usually link to external websites to encourage purchases, sign-ups, bookings, or the seeking of more information. The majority of CTAs we observed were ‘Learn More’, ‘Shop Now’, ‘Sign Up’, and ‘Send Message’ (Figure 17). While Facebook’s ad system includes a range of default CTA types, many of the ads that were captured used custom text. Nonetheless, the custom text often copied the default CTA types which might suggest that some page administrators are unaware of the default text options, or are opting for custom text for some other reason.




Figure 17. Count of Most Frequent CTAs (> 100 of this CTA type) used in ads in the Ad Observatory during 2022

Note: This visualisation utilises a log scale.



Advertisers

The most prominent advertisers in our dataset (see Figure 18) were from the following sectors:

- charity/social good (WWF Australia, Red Cross Lifeblood, The Smith Family, CHOICE)
- retail (The Oodie, Coles, Woolworths)
- media (Disney+, Paramount+)
- political (Australian Electoral Commission, Australian Unions, Getup!),
- insurance/banking/finance (Future Super, Health Insurance Comparison, Everyday Rewards)
- telecommunications (Telstra)

A number of highly observed advertisers also crossed or blended these sectors. For instance, Humble Bundle is a retail 'social good' enterprise which also sells digital game packages.

Most prominent pages in the Ad Observatory database




Figure 18. Facebook ads observed (aggregated by advertiser according to their page) in the Ad Observatory during 2022

Note: Colour indicates the number of unique (i.e. different) participants observing ads run by these pages, with red indicating more unique observers and green indicating fewer unique observers. For example ads from WWF Australia (dark red) were seen 2514 times by 423 different participants, while ads from Snag Tights (dark green) were seen 1069 times by 98 different participants.



Trends in Observation

Expanding the analysis of the number of users that observed advertising from specific advertisers, we note a linear trend relating the number of observations of ads (aggregated by Facebook page) with the number of unique users that observed at least one ad from one of these pages (see Figure 19). In other words, as ad volume by a specific Facebook page increases - which would likely be correlated with ad spend - so too the number of different participants that observe likely increases. A linear trendline equates:

$$\text{Number of observers} = (0.16 \times \text{number of observations}) + 2.15$$

This results in an R^2 value of 0.8, and p -value < 0.0001 . A p -value of < 0.05 is generally considered statistically significant and, as such, this result is statistically highly significant. Note that it would be impossible for any points to appear in the top left of the graph in Figure 19 as the count of observers can never be larger than the count of observations. Points that situate towards the bottom right of this graph are those ads which tend to be high repetition for a small subset of users. This equation suggests that ads are seen on average 6 times by the same user.

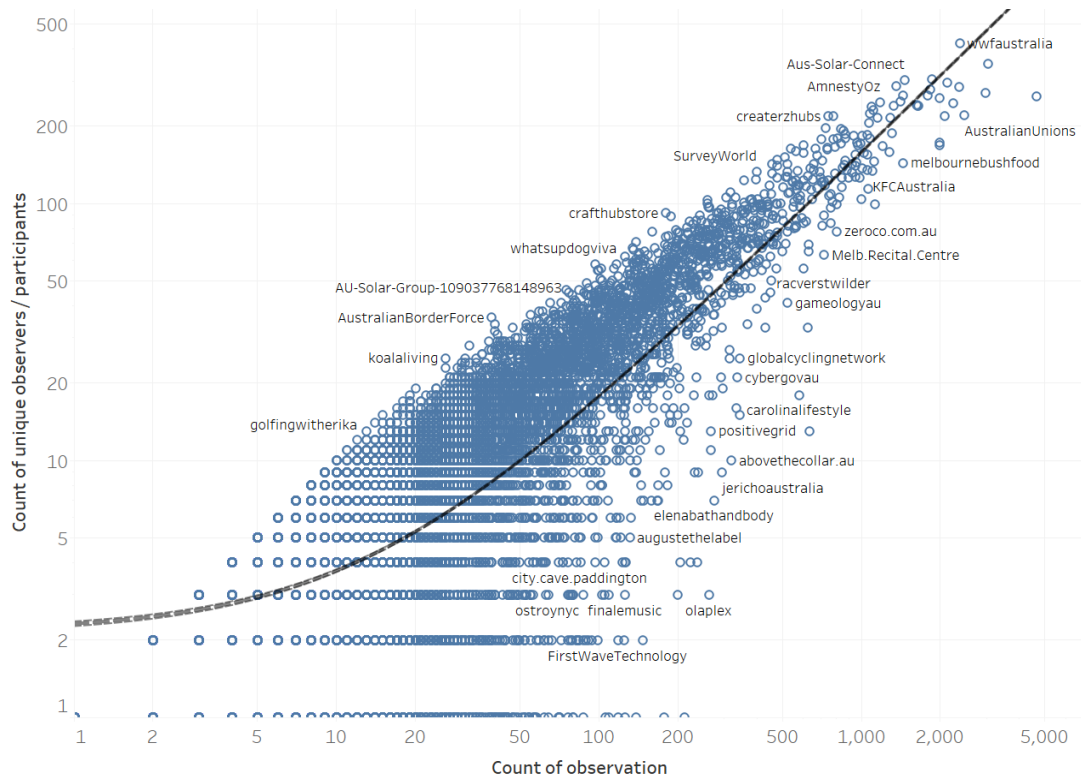


Figure 19. Scatterplot relating the count of ad observations to the count of unique observers/participants in the Ad Observatory during 2022

Note: The count of ad observations were aggregated by Facebook pages. A linear correlation between these values is also observed and is indicated using a trend line. This visualisation uses a log-log scale



5. WAIST (WHY AM I SEEING THIS AD) DATA

In response to public concerns and regulatory pressures, companies such as Meta (the parent company of Facebook) have introduced transparency mechanisms to inform people about the function of advertising on the platform. Such mechanisms include the Ad Library and the “Why Am I Seeing This Ad” (WAIST) feature. WAIST explanations appear as buttons within each ad, which take users to additional information about how the ad was targeted towards them, such as demographic or interest-based information (see Figure 20).

Despite being a central feature of Meta’s response to increasing external scrutiny, little is known about how the WAIST feature works or how it operates at a population level. The WAIST feature is fundamentally individualistic: it provides an explanation of how a single ad was served to a unique user at a particular time. There is no public information as to the distribution or prevalence of specific WAIST tags across advertising markets or user populations.

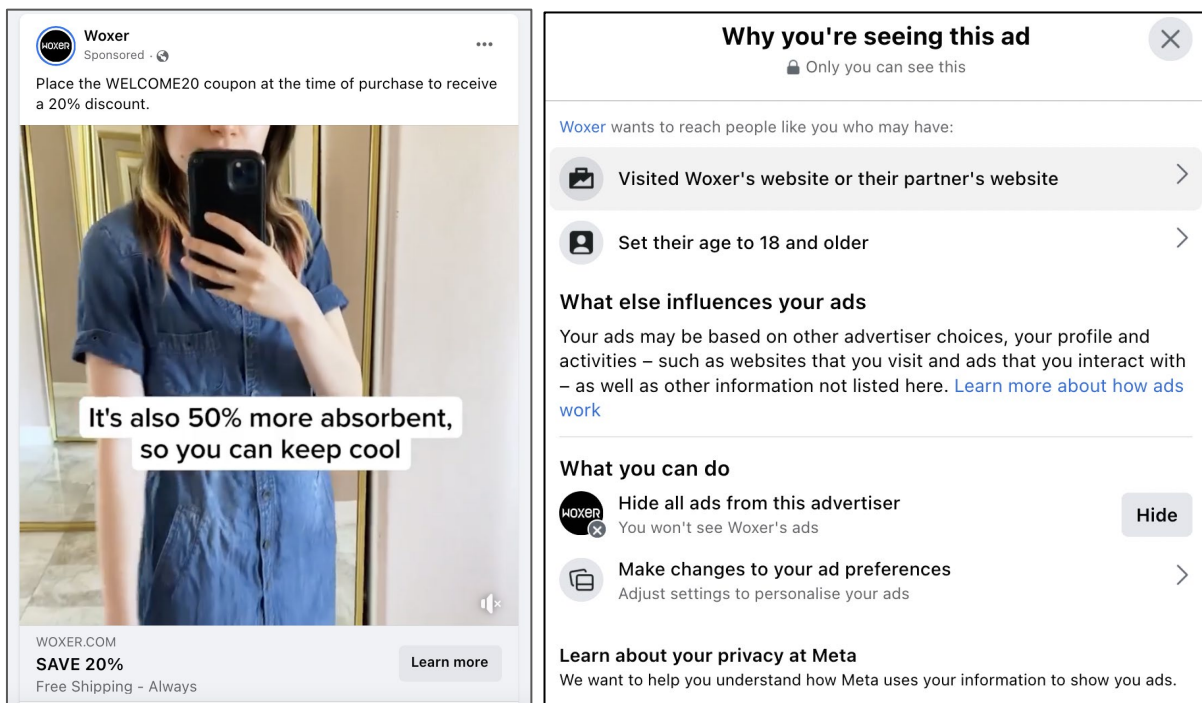


Figure 20. Example of “Why Am I Seeing This” (WAIST) explanation feature

Note: The user clicks on the three dots in the top right hand corner of the ad for the WAIST pop up to appear

By September 2023, the Ad Observatory had gathered more than 330,000 distinct ads and associated metadata, including WAIST categories of data, with 324,868 unique semantic labels within the WAIST ‘interest’ tag type gathered.



Our approach enables us to speculate on the alignment of WAIST data and patterns in our evidence about the platform’s actual ad targeting practices. Evidence of the platform’s ad targeting practices derives from advertisers’ use of the Facebook ad system. We are additionally able to speculate on the gaps between WAIST data and patterns in our evidence. Our ongoing analysis suggests that WAIST tags are unevenly deployed, with ‘AgeGender’ (including age ranges and binary gender classification) and ‘Location’ being the dominant WAIST tags applied to ads. WAIST ‘Interest’ tags are also prominent, while education, relationship, and employment status tags are rarely deployed (see Figure 21).

Number of Ads containing specific WAIST Type




Figure 21. Number of Ads Containing Specific WAIST Type

Relationships Among WAIST Tags

The WAIST Interest tag is a semantic description that relates to user interests, such as hobbies, popular culture, media, and other socio-cultural tastes and preferences. These interest tags can be considered as human-readable labels that map to more abstract numeric representations of user preference categories used by Facebook’s recommendation algorithms to associate and model users of the platform. These tags can organise into loose hierarchies, for example with tags like: ‘Music’, ‘Pop music’, or ‘Taylor Swift’. Interest tags also appear in sets, with many of these semantic tags appearing together within a single observation. For example, ‘Charity and causes’ may appear alongside ‘Community issues’. By mapping these relationships as a graph where each node represents an individual WAIST Interest tag, and edges are weighted according to the number of observations that use any two tags, it is possible to show the



topology of the WAIST interest code system and identify emergent patterns (see Figure 22). Among other things, this type of analysis can provide insight into the potential to reproduce protected or sensitive categories like race, sexuality, gender, and class based on proxy personalised ‘interests’ targeting, rather than targeting those categories explicitly (Phan & Wark, 2021).

The WAIST Interest tag is unique amongst other WAIST categories as it is the most open-ended, with 324,868 unique semantic tags discovered in our dataset, and likely many more in use by Facebook. Most other WAIST tags are simple Boolean or nominal sets without the same degree of diversity nor specificity. Interests are the third most utilised WAIST tag behind Age/Gender and Location.



Figure 22. Visual demonstration of WAIST Interest tag network, colours indicate clusters of closely related interest tags, organised around themes (music, food and beverage, sports, etc.)

Note: This network provides a unique and detailed perspective on the relationships among WAIST Interest tags within Facebook’s ad ecosystem, and highlights some of the multi-scale structure and emergent hierarchies of interest tags



6. RESEARCH QUESTIONS AND CASE STUDIES

6.1 Research Questions

Our work with the Ad Observatory data involves four main types of investigation and analysis:

1. Investigation of the volume and characteristics of online advertising. This advertising is often motivated by some question of broader policy significance, such as the prevalence of advertising for products and services that are known to cause harm (for instance, alcohol and gambling).
2. Analysis of the textual or symbolic qualities of different types of advertising. This work addresses policy-oriented questions, such as whether some advertisements appeal to harmful stereotypes or narratives. For instance, we consider how scam ads claim attention through the use of celebrities and a promise to secret knowledge, or how greenwashing ads make vague and unsubstantiated claims. Greenwashing refers to environmental claims that are vague, meaningless or unsubstantiated, meaning that consumers are given a misleading impression about the environmental credentials of a business or product (Rickards 2022).
3. Identification of how particular ads and collections of ads may be targeted at particular demographics such as gender, age, income or education, or at particular times. This search includes an examination of any patterns which suggest vulnerable groups may have been taken advantage of, and of any potential reflected or reinforced cultural patterns. The data collected can also provide insight into how demographic groups are targeted with particular types of ads, including in ways that may reproduce or exacerbate stereotyping and social sorting.
4. Mapping and exploring ad sequences that are 'tuned' to different people in terms of quality, nature, speed, and time (Lyons et al., 2022). We use a combination of statistical techniques and qualitative coding to identify clusters of individuals who receive particular types of ads - for instance, alcohol ads, gambling ads, and junk food ads - in patterns of frequency, time, and adjacency to other ads. We then identify common features of the ads themselves and the individuals receiving them.

6.2 Case Studies

The Ad Observatory research team is currently undertaking a number of case studies of pressing policy significance using the data collected by the Ad Observatory. In each case we



search the Ad Observatory dataset for relevant ads in order to conduct our case study analyses. Section 6 describes how we create each separate subset of data from the Ad Observatory to conduct our case studies. In conducting our case studies we focus primarily on types of advertising and/or advertising practices that could be considered harmful, either because they:

- promote products or services that often cause harm (such as alcohol and gambling);
- violate Australian consumer laws (such as consumer law protections against scams or prohibitions on deceptive and misleading advertising like greenwashing);
- are targeted at particularly vulnerable groups (such as ‘buy now pay later’ financial services or home delivery buttons in alcohol advertising) or at particularly vulnerable times;
- facilitate social sorting and stereotyping.


We are collaborating with consumer rights organisations, such as CHOICE, the Consumer Policy Research Centre (CPRC), and the Foundation for Alcohol Research and Education (FARE), in conducting some of these case studies. These will be published as joint reports and scholarly journal articles and publicised through the ADM+S website.

In the sections that follow we briefly describe six of our case studies: political advertising, financial advertising, scams, green claims, alcohol advertising, and gambling advertising. For the green claims, alcohol and gambling case studies we delve into more detail about the three different ways that a subset of ads was created in order to construct a dataset to support our case study work:

- the green claims case study was based on a thematic key word search to identify ads that contained misleading environmental claims;
- the alcohol case study was based on a ‘white list’ of alcohol brands and retailers who were the advertisers of alcohol products;
- the gambling case study combines both approaches.

6.3 Political Advertising

We begin our outline of our case studies with a brief description of our first investigation conducted with our partner, the ABC. In the first major use of the Ad Observatory, we uncovered an instance of political advertising that was in breach of Australian guidelines in regards to correct disclosure of party affiliation. During the lead up to the 2022 Australian Federal Election, an advertisement containing political logos of registered political parties (the



Australian Labor Party, the Liberal National Party, and the Australian Greens) was detected using the logo detection technology we developed. The advertisement referenced a group and a website that did not contain appropriate disclosures of party affiliation as required by the Australian Electoral Commission (AEC). An investigation by partner investigators at the ABC discovered that the website was linked to a One Nation Party candidate who had gained notoriety for their anti-science and anti-vaccine views, and who was running for election as a self-proclaimed 'freedom' candidate. Coverage by our project partners at the ABC ultimately led to the candidate being disendorsed by the One Nation Party and an investigation by the AEC (Bogle et al., 2022).

6.4 Financial⁵


We respond to concerns that financial products may be targeted at vulnerable grounds, including young people and people most at risk of financial hardship, by investigating ads that related to consumer financial products and services.

The research involves searching the Ad Observatory to identify how people are targeted with particular advertisements and collections of advertisements for financial products or services. We then interrogate the different types of products and services that are (or are not) being served to consumers online. Our research explores how financial consumerism and potential consumer harm is produced on social media and considers the implications for consumers' financial health and consumer protection. The research asks whether the targeting of financial ads appear to be related to particular personal characteristics or classes protected under anti-discrimination laws, such as financially vulnerable cohorts.

The research further compares the demographic distribution of ads against individuals and groups in our study population who have seen particular ads to understand what type of ads are being served by Facebook to these individuals and groups. We explore the distribution of ads along postcode areas (which we use here as proxies for race and ethnic affinity) and apply self-declared demographic details of participants to map how ads are served to audience segments and to establish if any harmful patterns are statistically significant.

Through this research, we seek to establish if particular demographic, protected groups are targeted with 'buy now pay later' ads, payday loans, and predatory financial services more broadly. Such targeting behaviour and some of these types of ads are prohibited by Meta's

⁵ The ADM+S researcher working on the financial case study is Kelly Lewis. She is working in partnership with CHOICE.



financial products and services advertising policy (Meta, n.d.). This approach also enables exploration of how the audience targeted by the advertising on Facebook compares with the demographic distribution of the eligible target audience selected by the advertiser, which provides a basis to examine whether targeting practices are compliant with the target market determination regulation made by the Australian Securities and Investments Commission (ASIC) as per its “design and distribution obligations” (ASIC 2020).

6.5 Scams⁶

We are investigating scam ads in the Ad Observatory, with a specific focus on ads which use unauthorized and ‘faked’ celebrity likenesses and endorsements to promote cryptocurrency and foreign exchange trading schemes.

The analysis of scam advertisements comes with a unique set of challenges. Unlike other advertisers, scam advertisers are not interested in building recognized brand identities that can be readily identified by platforms and critics. Our ad collection took place during a particularly active time for online scams. According to the Australian Competition and Consumer Commission (ACCC), Australians lost approximately \$3 billion to scams in 2022. This loss was up 80 per cent from the previous year. Many of these scams invite unwitting people to invest in dubious cryptocurrency or forex schemes. Scammers will often create convincing but fake web pages or apps to con people into passing over the details of their crypto wallet to thieves. Other scams include ‘rug pull’ crypto scams, where people are encouraged to invest in new types of crypto ‘coins’ only to have the scammer disappear once enough people have invested.

The ACCC and ASIC are trying to automatically detect and block scam websites. In 2023, they announced they had blocked 5,579 websites in three weeks, most of which were related to crypto scams. However, scammers are agile and hard to pin down. Often based in inaccessible overseas jurisdictions, these scams operate at such a scale that trying to block all of their IP addresses is a herculean task, even with automated tools. Many of these scams are advertised on social media, especially Facebook. The ACCC is currently suing Facebook’s parent company Meta for allowing ads which use unauthorised celebrity images to lend credibility to crypto ad scams (ACCC, 2022).

Recently, these ads have even included the unauthorized use of images of the Australian Prime Minister, Anthony Albanese, in addition to other well known media figures and celebrities. As

⁶ The ADM+S researchers working on scams are Mark Andrejevic and Chris O’Neill. We are working in partnership with CHOICE and the ABC on some of this work.



even a cursory inspection of the Facebook ad library reveals, the ACCC lawsuit does not seem to have had an impact on the prevalence of scam ads using celebrity images, nor on logos and bylines lifted from media outlets, including the ABC and major commercial news outlets in Australia.

Scam promoters have the advantage of not having to adhere to the truth or to ethical norms and guidelines of advertising. This means they can be more extreme and limited only by the willingness of the platform to respond to obvious policy violations. Rather than a straightforward endorsement, the ads we captured often evoke hidden knowledge. The ads also draw on familiar click-bait conventions, matching sensational headlines purporting forbidden knowledge with dramatically and obviously faked images of celebrities in apparently scandalous situations, such as being frogmarched in handcuffs or being bundled into the back of a police car.

By analysing the communication strategies and tropes of these ads in the Ad Observatory dataset, we can better understand the tactics that scammers are using to attempt to dupe Australian social media users and advise on heuristics and strategies to avoid becoming a victim of these operations.

6.6 Green Claims⁷

Many consumers increasingly want to make environmentally conscious purchases where they can, and many businesses are responding to this desire by promoting their environmental credentials. Recently, for example, the Australian Competition and Consumer Commission (ACCC) found that 57% of online businesses made ‘concerning claims’ about their environmental credentials with regard to cosmetics, clothing, and footwear. Food and drink sectors have the highest proportion of concerning claims (ACCC, 2023; Burry, 2022). These are claims that are vague, meaningless or difficult to substantiate on their face, meaning that they warrant investigation because they may amount to false and misleading advertising or ‘greenwashing’. Regulators in Europe and the UK have reported similar trends (Competition and Markets Authority, 2021).

This is concerning for consumers because those wanting to ‘make a difference’ with their purchases may be misled into buying products with inferior environmental performance. It is likewise concerning for businesses that are making progress in their environmental

⁷ The ADM+S researchers who are working on or have worked on green claims are Christine Parker, Loup Cellard, and Julian Bagnara. Some of our work on green claims is being done in partnership with CPRC.



performance but may still lose consumers to inferior environmental performing businesses. Misleading green advertising can also give the impression that production and consumption are less ecologically damaging and more environmentally sustainable than they really are. This could delay meaningful action on environmental problems, including the climate crisis. Such misleading conduct may contribute to a cultural environment that promotes the idea that sustainable consumption and business activities are already doing all that is needed to solve the climate crisis and other environmental predicaments. The complacency borne from this idea risks de-prioritising collective policy action, including stricter regulation.

More broadly, we expect to see green claims in advertising reflect, track, and possibly influence consumers' thinking about what counts as being 'green'. They also influence the perception of what business practices are seen as necessary and desirable by consumers as the climate and other environmental crises unfold. The Ad Observatory, as a large repository of a broad range of advertising, provides a unique window into green advertising practices and how Australians think about how green consumption and sustainable business is presented in a media environment often disconnected from natural environments.


6.6.1 Keyword Search Strategy

This case study used an iterative search of the Ad Observatory database for keywords that may likely return environmental or green claims. This was followed by a process of manual cleaning to identify irrelevant ads and manual coding of each ad by industry. Given the large size of the dataset generated (n = 55,600 observations), only a limited sample of the dataset was cleaned and coded which accounted for a majority (70%) of total observations.

Our aim was to create a single 'green claims' data set that was able to capture as many advertisements containing green claims as possible using a keyword search strategy. We achieved this in two stages: a pilot using six environmental keywords and then an expanded keyword search using 27 terms and Boolean operators to reduce noise.

The development of our list of search terms was based on two sources:

- Environmental terms identified by consumer regulators in Australia and elsewhere as likely to be problematic, such as environmentally friendly, sustainable, compostable, biodegradable, recycled, recyclable, and green (Rickard, 2022; ACCC, 2023).
- Responses to the CPRC's nationally representative survey of 2,000 Australians (conducted in October 2022) in answer to an open ended question: "What words and messages have you seen used to promote the green credential claims of a product or service?" (Burry, 2022).



The search included all advertisement text captured by the Ad Observatory, any metadata, and any additional text or metadata analysed from static image or video content through OCR or YOLOanalysis.

A search conducted on March 20, 2023 returned 19,707 unique advertisements by 5,430 unique advertisers. These were observed a total of 55,600 times by 926 unique participants in the Ad Observatory.

6.6.2 Cleaning and Coding

Next, we manually cleaned the data by deleting ads without a green claim. A broad and inclusive approach to determining 'green claims' was taken: any reference to green or environmental ideas and themes were included. This included non-commercial advertisements, such as ads made by government agencies and environmental advocacy organisations. Only completely irrelevant ads were discarded, which were those ads using the terms in a way that had nothing to do with being 'green' or 'the environment'. Figure 23 shows the ads by keyword and 'noise'. To ensure reliability, one researcher cleaned the data and another verified the cleaning. Ads were only cleaned out of the data set where both researchers agreed and inter-coder reliability was established. To ensure the policy relevance of our work, we also discussed with our advisors at the CPRC whether certain ad content should be considered as a green claim or not. At this stage, we also manually coded the ads by industry as the metadata sourced from Facebook did not provide a reliable or full account of the industry for each ad. The final dataset was used to produce a report, Seeing Green – prevalence of environmental claims on social media, in collaboration between the ADM+S and our partner Consumer Policy Research Centre in late 2023 (Gupte et al 2023, Parker 2023).



Keywords - Frequency and Noise




Figure 23. Environmental Keyword frequency and 'noise'

Because the extensive dataset comprised 19,707 distinctive ads from 5,430 distinct advertisers which were viewed 55,600 times, a complete cleaning and coding of the data was infeasible due to time and resource constraints. An initial evaluation of the data showed an approximate power law distribution with respect to the advertisers (see Figure 24), which indicated a significant proportion of ad observations were of ads made by a minority of advertisers. By focusing on these major advertisers, we optimised the cleaning and coding process while still representing a substantial portion of the data. The cutoff point was past the data saturation point (Lowe et al., 2018), which encompassed about 70% of the data observations (n = 38,962 out of 55,600 total observations). This was well past the inflection point of the pareto curve (Figure 25) and the point where any significant new insights could be gained, at least with respect to industry data and relevance.



The cleaned dataset includes 8,963 unique ads from 482 advertisers. These ads were observed 30,420 times by 832 separate individual participants. 1,800 ads constituting 8,494 observations were deemed irrelevant, constituting approximately 21.8% of the total 'noise'.

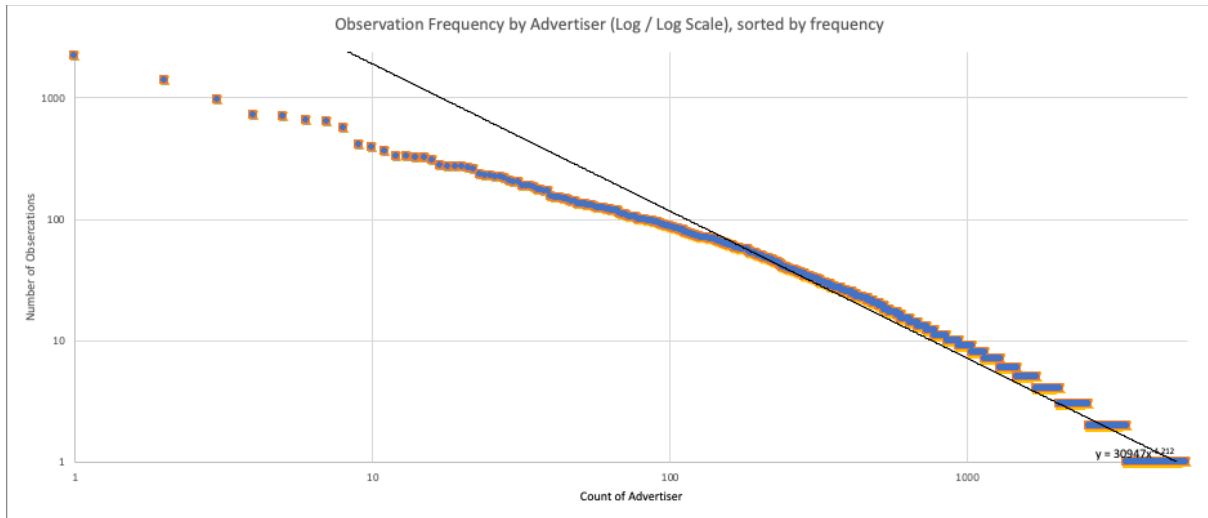


Figure 24. Observation of Ads with Green Claims (Frequency by Advertiser, Log Scale)

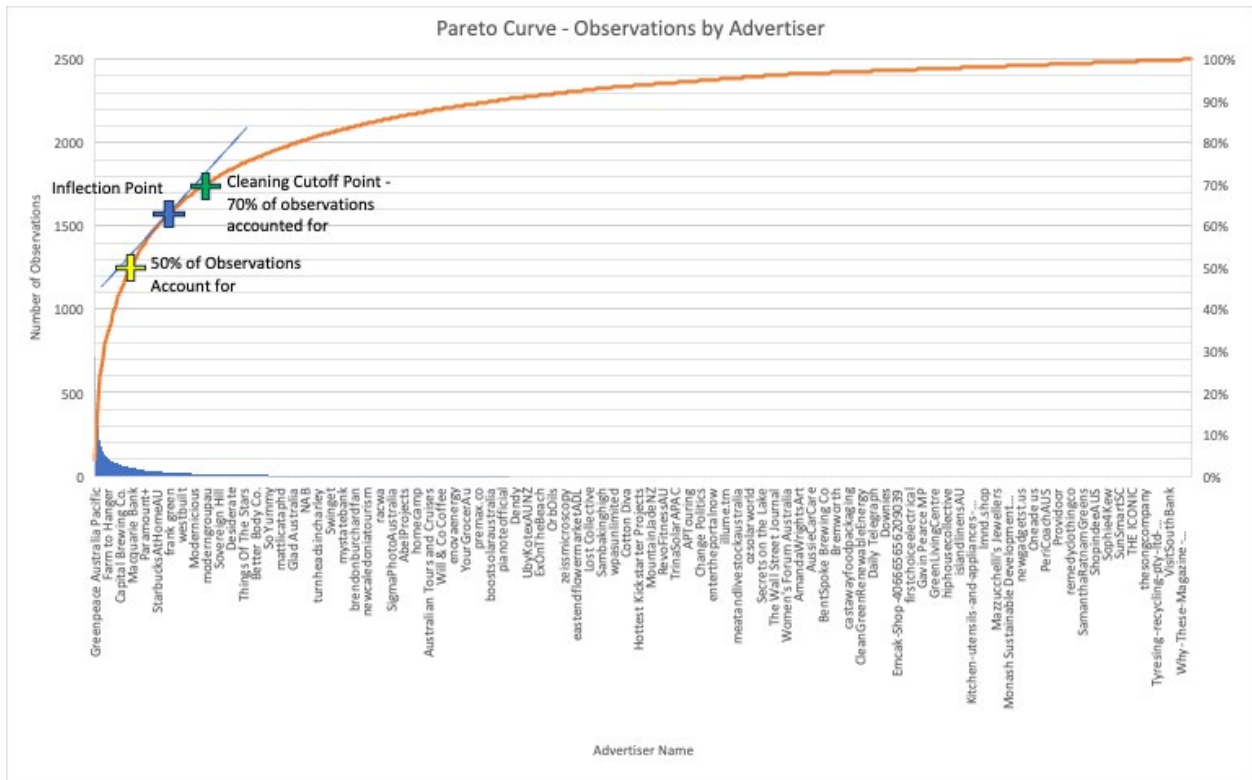


Figure 25. Pareto Curve cutting off the long tail of observations by advertiser



The dataset was further restricted to one calendar year (2022) for demographic analysis to ensure consistency with the timeframe of the descriptive statistics provided in the Participants and Demographics section above. The final dataset for these purposes consists of 8,498 unique advertisements from 466 advertisers resulting in 30,008 observations during this period.

6.7 Alcohol Advertising⁸

Alcohol advertising and its influence on drinking behaviour is a public health concern. Evidence shows that exposure to digital alcohol advertising increases the likelihood of consuming alcohol and promotes positive perceptions toward alcohol (Buchanan et al., 2018; Gupta et al., 2016). Alcohol advertising is strictly regulated in Australia for these reasons.

However, advertising on digital platforms poses new challenges in monitoring and evaluating the harms caused by alcohol consumption. First, digital advertising is largely unobservable to regulatory bodies, researchers, and civil society. While some efforts have been made to increase transparency, advertisements are only visible to the targeted users at one point in time which prohibits effective compliance monitoring (Hawker et al., 2022). Second, alcohol advertising on digital platforms is tuned towards conversion, meaning that ads are shown to individuals who are most likely to click (and ultimately purchase) the product advertised. This logic disproportionately impacts people who are at risk of harm from alcohol consumption based on their past behaviour.

The Ad Observatory addresses both concerns by providing a real-time view into when and where alcohol advertising appears in the newsfeeds of Australian Facebook users. Our alcohol case study explores the nature of alcohol advertising on Facebook and how chronological sequences of ads can offer a new perspective on patterns of harmful advertising.

6.7.1 Compiling the List of Advertisers

For this case study, we compiled a dataset of ads based on a list of known advertisers' Page IDs. This method was chosen for alcohol advertising because the production and distribution of alcohol in Australia is dominated by a few major corporations. We first developed a spreadsheet to collect the name of key advertisers, the URL to their Facebook page, their Facebook Page ID

⁸ The ADM+S researchers working on the alcohol case study are Nic Carah and Lauren Hayden. Some of our work on alcohol advertising is being done in partnership with FARE.



number, and their market category based on whether the advertiser was an alcohol producer, retailer, or venue.


Given the relatively concentrated alcohol industry in Australia, major alcohol companies were identified via an industry report (Beck et al., 2023). This report included producers of beer, wine, spirits, and ready-to-drink products (RTDs) as well as national retailers. All consumer-facing brands listed in the report were added to the spreadsheet. For example, Budweiser, Corona, and Stella Artois were added to the list but the holding company (Anheuser-Busch InBev) was not. To ensure that the maximum number of advertisers were collected, we reviewed each holding company website for any additional product brands that were not listed in the report. This resulted in 354 alcohol products and retail brands. To capture brands not owned by a major alcohol company (such as craft beers), we copied a product list from the Dan Murphy's website. In order to make this task manageable, only craft beers and RTDs were collected from this catalogue, as these products were identified as the products of interest in terms of advertising on Facebook. This totaled 259 advertisers included in our analysis.

6.7.2 Identifying Facebook Page ID Based on URL

Using the URLs for the Facebook page, we were able to computationally collect the unique Page IDs for each advertiser. The Page ID acts as an identifying variable within the Ad Observatory data set to flag alcohol advertisers.

6.7.3 Join List of Advertiser IDs to Australian Ad Observatory Database in Tableau for Analysis

The spreadsheet of alcohol advertisers was joined to the Australian Ad Observatory cloud database in Tableau. The listed Page IDs were then used to isolate observations of the identified advertisers for analysis. This method allowed use of the spreadsheet variables, such as market categories, to analyse the joined up dataset. Four additional alcohol advertisers were identified throughout the analysis process and later added to the list for a total of 263 alcohol producers and retailers in our dataset. For the period from 6 December 2021 to 6 December 2022, we identified 2,624 distinct alcohol ads and 7,568 distinct observations.



6.8 Gambling Advertising⁹

Australia is the world leader in gambling losses (Australian Gambling Research Centre 2023). According to the Australian Gambling Research Centre, 73% of Australians gamble at least once in a twelve month period (2023). Beyond being a commonly-reported nuisance by members of the public, gambling advertising carries risks for people with compulsive gambling problems. Illegal gambling flouts the gambling protection measures in place to combat issues with gambling addiction, such as the BetStop self-exclusion measure. Likewise, gambling ads on Facebook that fail to comply with Australian laws restricting gambling advertising expose problem gamblers to real issues. These issues should be addressed in order to protect at-risk Australians and their families.

In an investigation conducted in partnership with the ABC as part of the gambling advertising case study, we uncovered gambling advertising by Bitstarz, an online offshore casino that cannot legally operate nor advertise in Australia (see Figure 26). The Ad Observatory captured four instances of advertisements on Facebook that appeared to have been targeted at Australians. Their images included graphics of Australian flags and metadata indicating that the person being served the ad was located in Australia. Not only were these ads in circumvention of local requirements, they appeared to be in violation of Meta's own policies which seek to prevent the hosting of illegal gambling advertising in countries where such advertisements are prohibited. This finding was reported on national broadcast news and current affairs programs on the ABC (Briggs, 2023; Parker et al., 2023).

6.8.1 Search Strategy: Using Regulatory Processes to Filter Gambling Data

Mapping the experiences of gambling is a challenging process due to the substantial amount of illicit and gray zone gambling materials in circulation and the degree to which gambling is embedded in Australian social life. Like alcohol and alcohol advertising, gambling and gambling ads are regulated in Australia, which allow us to track specific known providers. While some ads from legal gambling providers do not use much gambling language, the use of business identifiers makes it easier to track regulated gambling companies. However, deceptive, unregulated, and illegal gambling may also be advertised online in significant quantities, including on Meta's platforms.

⁹ The ADM+S researchers working on the gambling advertising are Mark Andrejevic, Cesar Albarran-Torres, and Robbie Fordyce. Research assistance has been provided by Sijun Shen, Jenn Wilson, and Isabella Mahoney.




Figure 26. Images of Bitstarz advertising from ABC reporting

Note: From Briggs (2023) . Images are pixelated to remove promotional messages.

A further complication is the degree to which gambling-like behaviour is present across Australian society. Gambling and gambling terminology are present across ads for sports, charities, gaming, apps, and astrology businesses. Such ads include charities advertising a raffle for fundraising and a horoscope ad suggesting a big lottery win. Some gambling operators may also pose as charities to avoid detection while providing gambling services, making it even more difficult to identify gambling ads. As such, identifying gambling ads is a 'fuzzy problem' that cannot be neatly solved by searching either advertisers or keywords in isolation. Addressing the problem requires a combination of both.

6.8.2 Initial Pilot Search

Initially, we developed a wide-net search of gambling content through the Ad Observatory using general gambling terminology that involved a mix of gambling practices ('slots', 'pokies'), gambling outcomes ('jackpot', 'win'), and gambling organisations ('Sportsbet', 'casino'). This



search identified 42,693 ads that used gambling language, as well as an expected degree of 'noise'. The majority of these ads met our criteria of gambling as "the commercialisation of controlled experiences of risk" (Parker et al.,2023).

However, this search also produced many results that were unlikely to meet conventional understandings of gambling. Nonetheless, such a permissive dataset is informative about the extent of gambling terminology and metaphors existing throughout Australian advertising and might inform a research perspective on gambling culture. A permissive dataset would not be as useful in addressing people's experiences of gambling advertising as more narrowly understood. We therefore sought to develop a tighter dataset through a combination of more precise targeting relying on regulatory materials.

6.8.3 Advertiser Search

Operators of gambling services are required to be registered with authorities at the Australian state or territory level. Directories are available for registered operators. These registers were useful in constructing our search as they allow us to identify advertisers in the dataset by name or through other identifiers, such as their embedded web address or persistent identifier key.

Searching by compiled key terms of advertisers assisted us in producing a dataset of gambling companies that are registered within Australia. This dataset allowed us to gain insight into the rate at which the dominant gambling companies advertise on Facebook. Within this dataset, however, we found advertisers masquerading as legally-compliant gambling providers, in some cases using the branding of legal operators in their own ads. The Ad Observatory is capable of identifying these ads even though they cannot be readily found using Meta's own Ad Library tool.

6.8.4 Keyword Search

Searching for regulated advertisers usefully exposes some grey market gambling ads. However, illegal gambling ads are not detectable through searches based on advertiser names. For example, in some cases Facebook allows advertisers providing gambling products or services from overseas to advertise on their site without registration and without adherence to Australian law (Parker et al.,2023). The keyword search process was therefore developed to identify gambling advertising that was reaching an Australian context but was not otherwise picked up through our other methods of searching by advertiser. The process is especially useful as some advertisers changed their page names or only released a small batch of targeted ads (with many of these ads having since been removed from Meta's Ad Library).



The keyword system was difficult to refine due to the tendency for shorthand terms in gambling ads, such as ‘TAB’ and ‘bet’, and other Australian phrases such as ‘FS’ (free spin) and ‘multi’. This complicated the searching process. Due to these restrictions we focused only on several key metrics in the overall dataset. For gambling ads, we limited our search to the label, title, text, and call to action. For the advertiser, we limited our search to the advertiser name and embedded web links. This helped to control the issues related to coincidental search matches on sentence fragments that were unrelated to gambling, such as fragments related to WAIST categories and other internal data. Given the limits of the keyword search process, we used a search scoring system employed via fine-tuned text searches. We incorporated the summative score for each ad into our final dataset.

6.8.5 Classification of Gambling Products

Because of its significance within Australia, gambling is well-researched across many disciplines, agencies, and NGOs. The Queensland Government’s Australian Gambling Statistics project has publicly released data on Australian gambling habits since 1993. The City of Dandenong in Victoria, which has an acute problem gambling issue, has substantial behavioural and statistical data about gambling within local government areas in Victoria. These data provide insights into pathways leading to problem gambling.

Other research from government bodies, such as the Australian Communication and Media Authority and BetterHelp, and from NGOs such as GamblingHelpOnline and GamblersHelp, provided us with useful frameworks for segmenting our collected ads into different gambling categories. Based on the classifications from these organisations, combined with additional categories drawn from our analysis, we have identified the following categories of gambling ads.

Table 1. Description of the categories of gambling ads in Australia

Category	Description
Bookmaking	Outcome based betting, not otherwise mentioned.
Casino	A dedicated gambling location.
Gaming	Specific casino-based games, including blackjack, roulette, poker, craps, and other similar games.



Culture	Ad for a gambling cultural group rather than product, such as social hubs, sites, or similar group where no specific product is centrally advertised.
Keno	Ad for a numbers game, including Lotto and Powerball.
Lottery	Ad for a random-draw game, including raffles.
Online	Gambling that is accessed online. Distinct from social gambling due to the possibility of payout.
Pokies	Electronic Gaming Machines (EGMs), commonly referred to as 'pokies' in Australia.
Racing	Outcome based gambling, including gambling on dog or horse racing.
Scratchies	A scratch-off ticket.
Social	Simulated online gambling. Distinct from online gambling due to the lack of direct monetary payout.
Sport	Betting on the outcome or progress of competitive sports. Commonly known as 'sports betting'.

Additionally we created further categories based on key secondary factors.

Table 2. Description of Types of Gambling Ads and Advertisers in Australia

Category	Description
Anti-gambling	Ad from an NGO or government designed to reduce gambling prevalence.
Registered	An advertiser of gambling products who is registered in an Australian state or territory.
Unregistered	An advertiser of gambling products who is not registered in Australia.



Fundraising	Advertises gambling as a fundraising activity.
Deceptive	Ad uses branding or materials to conceal its origins, operations, or purpose. May or may not be legal.

6.8.6 Reviewing Offending Ads

The origin of gambling-related ads within Australia can be identified with confidence using a combination of regulatory tools available to the public. As with almost all ads, gambling ads link to an external website. For Australian gambling ads, most use the local top-level-domain “.com.au”. This domain requires registration with an Australian Business Number (ABN), which is present in domain lookup requests via Australian ISP ‘whois’ requests. The ABN can be cross-referenced with information on state level registered gambling providers or against the Australian Charities and Not-For-Profit Register. This makes the process of identifying a legally registered advertiser much easier.

7. ETHICAL AND COMPLIANCE CONSIDERATIONS



The research team has considered various ethical and compliance issues that arise when working with and publishing from the Ad Observatory, These issues include copyright, defamation, and liability issues that may arise when re-publishing misleading or illegal advertisements, as well as our ethical obligations to research participants and others involved in the research.

Ethics and privacy

Our main obligation to our participants is to ensure that there is no way in which they could be reidentified by the data we publish. While they are anonymous in all the data we collect, it may be possible in some situations that individuals could be re-identified by a combination of otherwise anonymous data. We are therefore careful to ensure that we do not publish anything that could re-identify any individual who contributed to our dataset. We are also observant of the privacy of any individuals who may be advertisers in the ads we collect.

Copyright

We rely on the exception to copyright in relation to 'fair dealing for research and study purposes' in s40 of the *Copyright Act 1968* (Cth). There are public interest purposes for



conducting our research project. Supporting factors for the use of this exception include the fact that the images we captured are not highly artistic works which creators might be deeply and personally connected to. Further, it would be unreasonably burdensome for our research to obtain licences from each advertiser for every advertisement we collect.

Defamation and Allegations of Illegality or Misconduct

We are careful to make sure that any conclusions we draw about misleading or illegal conduct are backed up by evidence. We are mindful of the risks of defamation against individuals and small businesses.

Re-publishing Illegal or Misleading Content in Ads

We are careful to make sure it is clear from the context of our papers and presentations that we are not in any way endorsing or adopting as true any advertisements that we discuss or reproduce. In particular where ads are clearly illegal or harmful, such as scam ads, some gambling ads, and ads containing misinformation about health and safety issues, we strive to caption the ad clearly and prominently indicate that we are not endorsing it. In some cases we watermark the ad, depending on where the material is published and to which audience.



8. REFERENCES

- Andrejevic, M., Fordyce, R., Li, N., Trott, V., Angus, D. & Tan, J. (2021). Unregulated and segmented dark ads on social media: Consumer education and regulatory options. Australian Communications Consumer Action Network. <https://apo.org.au/node/313053>
- Araujo, T., Ausloos, J., van Atteveldt, W., Loecherbach, F., et al. (2022). OSD2F: An open-source data donation framework. *Computational Communication Research*, 4(2), 372-387.
- Australian Bureau of Statistics. (2021). 2021 Census of population and housing: General community profile. Australian Bureau of Statistics. <https://www.abs.gov.au/census/find-census-data/community-profiles/2021/AUS>
- Australian Competition and Consumer Commission. (2022). ACCC takes action over alleged misleading conduct by Meta for publishing scam celebrity crypto ads on Facebook. *Media Release 34/22*, 18 March. <https://www.accc.gov.au/media-release/accc-takes-action-over-alleged-misleading-conduct-by-meta-for-publishing-scam-celebrity-crypto-ads-on-facebook#:~:text=The%20ACCC%20alleges%20that%20the,and%20former%20NSW%20Premier%20>
- Australian Competition and Consumer Commission. (2023). Greenwashing by businesses in Australia: Findings of the ACCC's internet sweep of environmental claims. <https://www.accc.gov.au/system/files/Greenwashing%20by%20businesses%20in%20Australia.pdf>
- Australian Electoral Commission. (2022). 2022 Federal Election House of Representatives First Preferences by Party [Data Set]. <https://results.aec.gov.au/27966/Website/HouseStateFirstPrefsByParty-27966-NAT.htm>
- Australian Gambling Research Centre. (2023). Gambling participation, experience of harm and community views: An overview. Australian Institute of Family Studies. https://aifs.gov.au/sites/default/files/2023-03/2302-overview_gambling-participation-harm-views.pdf



Australian Securities and Investments Commission (ASIC). 2020. *Regulatory Guide 274: Product Design and Distribution Obligations*. Commonwealth of Australia.

<https://asic.gov.au/regulatory-resources/find-a-document/regulatory-guides/rg-274-product-design-and-distribution-obligations/>

Allf, B., Cooper, C., Larson, L., Dunn, R., Futch, S., Sharova, M. and Cavalier, D. (2022). Citizen science as an ecosystem of engagement: implications for learning and broadening participation. *BioScience* 72(7): 651–663.

Briggs, C. (2023). Online casinos based offshore are illegally targeting Australians on Facebook. Who is responsible? *ABC News Online*, 21 March. <https://www.abc.net.au/news/2023-03-21/illegal-online-gambling-facebook-meta-casino/102118408>

Beck, E., Pierce, H., Stafford, J. & Wilson, C. (2023). *A Guide to the Alcohol Industry, Major Alcohol Companies in Australia: Producers and Distributors*. Perth, WA: Cancer Council Western Australia.

Bogle, A., Karagic, D. & Nguyen, K. (2022). Pauline Hanson's One Nation asked to please explain why some of her candidates are invisible in the seats they're supposed to be contesting. *ABC News Online*, 5 May. <https://www.abc.net.au/news/2022-05-05/pauline-hanson-one-nation-asked-to-explain-invisible-candidates/101039240>

Buchanan, L., Kelly, B., Yeatman, H. & Kariippanon, K. (2018). The Effects of Digital Marketing of Unhealthy Commodities on Young People: A Systematic Review. *Nutrients* 10(2): 148.


Burgess, J., Andrejevic, M., Angus, D. & Obeid, A. (2022). Australian Ad Observatory: background paper. *ARC Centre of Excellence for Automated Decision-Making and Society*, 15 July. <https://apo.org.au/node/318616>

Burry, K. (2022) The consumer experience of green claims in Australia. *Consumer Policy Research Centre*, December. https://cprc.org.au/wp-content/uploads/2022/12/CPRC-Green-Claims_Final.pdf

Competition and Markets Authority. (2021). Global sweep finds 40% of firms' green claims could be misleading [Press release]. *GOV.UK*, 28 January. <https://www.gov.uk/government/news/global-sweep-finds-40-of-firms-green-claims-could-be-misleading>



- Gupta, C., Bagnara, J., Parker, C. and Obeid, A.K. (2023). *Seeing Green – prevalence of environmental claims on social media*. Consumer Policy Research Centre (CPRC) and ADM+S. <https://cprc.org.au/seeing-green/>
- Gupta, H., Pettigrew, S., Lam, T. & Tait, R. J. (2016). A systematic review of the impact of exposure to Internet-based alcohol-related content on young people's alcohol use behaviours. *Alcohol and Alcoholism (Oxford, Oxfordshire)* 51(6): 763–771.
- Hawker, K., Carah, N., Angus, D., Brownbill, A., Tan, X. Y., Dobson, A. & Robards, B. (2022) Advertisements on digital platforms: How transparent and observable are they? *Foundation for Alcohol Research and Education*, 2 September. <https://fare.org.au/transparency-report/>
- Hughes, C. (2023) Australia: Facebook user distribution by age 2022. *Statista*, 24 May <https://www.statista.com/statistics/680607/australia-facebook-user-distribution-by-age-and-gender/>
- Lowe, A., Norris, A., Farris, A. J. & Babbage, D. (2018). Quantifying Thematic Saturation in Qualitative Data Analysis. *Field Methods* 30(3): 191-207.
- Lyons, A. C., Goodwin, I., Carah, N., Young, J., Barnes, A. M. & McCreanor, T. (2023). Limbic Platform Capitalism: Understanding the Contemporary Marketing of Health-Demoting Products on Social Media. *Addiction Research & Theory* 31(3): 1–6.
- Meta, inc. (undated). Prohibited financial products and services. https://transparency.fb.com/en-gb/policies/ad-standards/deceptive-content/prohibited-financial-products-and-services/?source=https%3A%2F%2Fwww.facebook.com%2Fpolicies_center%2Fads%2Fprohibited_content%2Fprohibited_financial_products_and_services%3Fwtsid%3Drdr_0RvA0ANUGaTOKsQwb
- NYU Cybersecurity for Democracy. (2024) *NYU Ad Observatory* at <https://adobservatory.org/>
- Park, S., McGuinness, K., Fisher, C., Lee, J. Y., McCallum, K. & Nolan, D. (2022). *Digital News Report: Australia 2022*. Canberra: News & Media Research Centre, University of Canberra.



- Parker, C. (2023). Social media ads are littered with green claims. How are we supposed to know they are true? *The Conversation*. 1 December. <https://theconversation.com/social-media-ads-are-littered-with-green-claims-how-are-we-supposed-to-know-theyre-true-218922>
- Parker, C., Albarrán-Torres, C., Briggs, C., Burgess, J., Carah, N., Andrejevic, M., Angus, D., & Obeid, A. (2023). Addressing the accountability gap: gambling advertising and social media platform responsibilities. *Addiction Research and Theory*. DOI: [10.1080/16066359.2023.2269852](https://doi.org/10.1080/16066359.2023.2269852)
- Phan, T. & Wark, S. (2021). What personalisation can do for you! Or: how to do racial discrimination without 'race'. *Culture machine* 20: 1-29.
- ProPublica. (2020). Facebook Political Ad Collector. *ProPublica*, 16 April. <https://projects.propublica.org/facebook-ads/>
- Queensland Government Statistician's Office. (2022). Australian Gambling Statistics 37th edition, 1994–95 to 2019–20. *Queensland Treasury, December*. <https://www.qgso.qld.gov.au/issues/2646/australian-gambling-statistics-37th-edn-1994-95-2019-20.pdf>
- Rickard, D. (2022). Speech to Sydney Morning Herald Sustainability Summit. *Australian Competition and Consumer Commission*, 20 September. <https://www.accc.gov.au/speech/speech-to-smh-sustainability-summit>



admscentre.org.au



Australian Government
Australian Research Council

This Centre is funded by the
Australian Government
through the Australian
Research Council

