

# PRE-ANALYSIS PLAN

*AI-Generated Political Information and Democracy*

**Authors:**

*Professor Nicholas Biddle (Australia National University)*

*Dr Svitlana Chernykh (Australia National University)*

*Dr Constanza Sanhueza Petrarca (Australia National University)<sup>1</sup>*

Version: December 11<sup>th</sup>, 2025

## Summary

This pre-analysis plan outlines the research design and analytical strategy for a vignette survey experiment examining how the use of AI affects public perceptions of election integrity, trust in investigating authorities, confidence in elections and support for democracy. The experiment is embedded within the December 2025 ANUpoll conducted online with approximately 3,500 adult Australian respondents.

---

<sup>1</sup> Corresponding author: [Constanza.sanhueza@anu.edu.au](mailto:Constanza.sanhueza@anu.edu.au)

## Table of Contents

1. INTRODUCTION .....	3
2. RESEARCH QUESTIONS AND HYPOTHESES.....	4
2.1 Main Effects.....	4
2.2 Interaction Effects.....	5
3. EXPERIMENTAL DESIGN .....	5
4. OUTCOME MEASURES.....	6
5. FIELDWORK .....	6
5.1 Survey .....	6
5.2 Factorial Analysis (3x3x3) .....	6
5.3 Statistical Power .....	6
6. ANALYSIS.....	7
6.1 Primary Analysis .....	7
6.2 Effect Size Interpretation and Substantive Significance.....	7
6.3 Subgroup Analyses .....	8
6.4 Robustness Checks .....	8
7. ETHICS.....	8
8. TIMELINE AND REPORTING.....	8
9. SURVEY QUESTIONS DECEMBER 2025.....	9
A. AI AND DEMOCRACY .....	9
B. REGULATIONS.....	11
C. VIGNETTE EXPERIMENT .....	13

## **1. INTRODUCTION**

Artificial intelligence (AI) is increasingly shaping democracy, with significant implications for elections and governance. As generative AI technologies become more integrated into political communication and campaign strategies, concerns have emerged regarding their potential to undermine voters' perceptions, public trust, manipulate electoral outcomes, and erode the foundations of democratic representation. While governments around the world grapple with how to regulate these emerging technologies, public concern is palpable and growing.

Generative AI technologies, particularly large language models like ChatGPT, possess unprecedented capabilities to produce human-like content at scale—including text, images, and videos—that can be weaponized to create disinformation, deepfakes, and other forms of manipulated media. Such content has substantial impacts on the digital public sphere, threatening democratic representation, accountability, and citizen trust in political institutions (Feldstein 2019; Kreps and Kriner 2023; Coeckelbergh 2025). As Kreps and Kriner (2023) describe, the most threatening aspect of generative AI is the speed and volume at which it can produce content, potentially drowning political communication channels and making it increasingly difficult for citizens to distinguish between genuine information and AI-generated manipulation. This fundamentally hampers the quality of democratic representation, as both political actors and citizens struggle to navigate a polluted information ecosystem where authentic public opinion becomes indistinguishable from algorithmically manufactured consensus.

Despite growing recognition of these risks, empirical research on public perceptions of AI threats to democracy remains limited, particularly in the context of actual electoral cycles. While recent scholarship has begun to explore AI-powered threats from theoretical and ethical perspectives (Jungherr 2023; Kreps and Kriner 2023; Coeckelbergh 2025), there is a critical need for systematic empirical analysis of how citizens perceive these threats, whether these perceptions change over time as they are exposed to AI-generated content during campaigns, and how regulatory interventions might shape public attitudes toward AI governance in democratic contexts. Understanding public opinion on AI regulation is crucial for developing effective policy responses that balance innovation with democratic protection, as democratic legitimacy requires that governance frameworks reflect citizen preferences and concerns.

This study addresses this gap by analysing panel survey data collected before and after the 2025 Australian federal election, examining how Australians perceive AI threats to democracy, their trust in electoral processes potentially affected by AI-generated content, and their preferences for regulatory interventions. Through a combination of observational panel analysis and an embedded survey experiment, we investigate both the evolution of attitudes toward AI regulation over the course of an election campaign and the causal effects of exposure to hypothetical AI-threat scenarios on democratic attitudes. Our research contributes to the growing literature on AI and democracy by providing the first empirical evidence of how real-world electoral experiences with AI-generated content shape public perceptions of democratic integrity and regulatory preferences.

## 2. RESEARCH QUESTIONS AND HYPOTHESES

This study seeks to investigate the following research questions:

1. How and under which circumstances does AI-generated political information affect citizens' perceptions of elections and democracy?
2. What can be done to mitigate the negative impact of AI on democracy?

The experimental design presented in this pre-analysis plan manipulates three key dimensions associated with AI to assess their effects on democratic perceptions: (1) the investigating authority, (2) the prevalence of AI-generated false information, and (3) voter ability to identify AI-generated political content. Based on theoretical expectations about institutional trust, information processing, and democratic resilience, we will investigate the following working hypotheses:

### 2.1 Main Effects

#### H1: Investigating Authority Effect

The investigating authority will influence trust in findings:

- H1a: *Electoral Commissions investigations receive highest trust (official government authority)*
- H1b: *Fact-checking organizations investigations receive high trust (perceived independence)*
- H1c: *Social media platforms investigations receive lower trust (conflict of interest)*

The investigating authority will influence trust in elections and democracy

- H1d: *Electoral Commissions investigations result highest trust in elections and democracy*
- H1e: *Fact-checking organizations investigations result highest trust in elections and democracy*
- H1f: *Social media platforms investigations do not affect trust in elections and democracy*

#### H2: AI-Content Prevalence Effect

Higher prevalence of AI-generated false information will lead to:

- H2a: *Greater perceived threat to election integrity*
- H2b: *Lower confidence in election legitimacy*
- H2c: *Higher likelihood of believing AI influenced the outcome*
- H2d: *Lower satisfaction with democracy*
- H2e: *Greater concern about similar situations occurring domestically*

#### H3: Voter Ability Effect

Higher voter ability to identify AI content will lead to:

- H3a: *Lower perceived threat (voters can protect themselves)*
- H3b: *Higher confidence in election legitimacy*
- H3c: *Lower likelihood that AI influenced the outcome*
- H3d: *Higher satisfaction with democracy*

## 2.2 Interaction Effects

**H4: Investigating Authority × Prevalence:** Trust in the investigation may vary depending on the severity of the problem found.

**H5: Prevalence × Voter Ability:** Threat perception will be moderated by voter AI-identification ability.

## 3. EXPERIMENTAL DESIGN

The experiment employs a 3×3×3 factorial between-subjects design with the following manipulations:

**Table 1: Vignette Experiment Dimensions and Levels**

<b>Dimension 1 Investigating Authority</b>	<b>Dimension 2 Prevalence of AI-generated False Information</b>	<b>Dimension 3 Voter Identification Ability</b>
1: Electoral Commission	1 (High): dominated online content related to the election.	1 (Low): Almost no voters could identify AI-generated political content
2: Independent fact-checking organizations	2 (Moderate): was found in a moderate amount of online content related to the election.	2 (Moderate): Some voters could identify AI-generated political content.
3: A Social media Platform	3 (Low): was only minimally found in online content related to the election.	3 (High): Most voters could identify AI-generated political content.

Respondents will be randomly assigned to one of the 27 experimental conditions (3×3×3). The vignette will describe a hypothetical scenario in which the investigating authority examined the use of AI-generated content during an election campaign and reached conclusions about its prevalence and voters' ability to identify such content.

**Table 2: Hypothetical Scenario Use of AI in Federal Election**

### **SCENARIO:**

We want to ask you about a hypothetical situation about the use of AI in election campaigns.

*Imagine that a federal election took place in a hypothetical country. After the election, [INVESTIGATING AUTHORITY] conducted an investigation into the use of AI to generate political content during the election campaign. The investigation found that AI generated false information [PREVALENCE LEVEL]. Additionally, the investigation revealed that [VOTER IDENTIFICATION LEVEL].*

## 4. OUTCOME MEASURES

- Satisfaction with democracy (C 18)
- Trust in findings about AI-generated political content (C19)
- Trust in election outcomes (C20)
- Trust in legitimacy of election (C21)
- Regulatory preferences (C22)
- Concerns about similar situation in Australia (C23)

## 5. FIELDWORK

### 5.1 Survey

The experiment is embedded within the December 2025 ANUpoll. The particular wave of the survey series will be conducted on the Online Research Unit's (ORU's) Australian Consumer Panel. The survey commenced with pilot data collection of around 70 respondents on the 9<sup>th</sup> of December, 2025. It is expected that data collection will be completed by mid-December, with an eventual sample size of around 3,500 respondents.

The target sample is 1,000 respondents aged 18 to 24 years, and the remaining 2,500 respondents aged 25 years and over.

Survey weights will be used in the analysis, calculated using Iterative Proportional Fitting (IPF), or raking. Population benchmarks that will be used for weighting purposes are age, sex, education, and current employment. The first two of these measures come from population estimates from the Australian Bureau of Statistics, the third (education) from the 2021 Census, and the fourth (employment) from the September 2024 Labour Force Survey. Bootstrap Standard errors will be estimated in STATA using 250 replications and seed 10121978.

The data will be made available for secondary analysis through the Australian Data Archive.

### 5.2 Factorial Analysis (3x3x3)

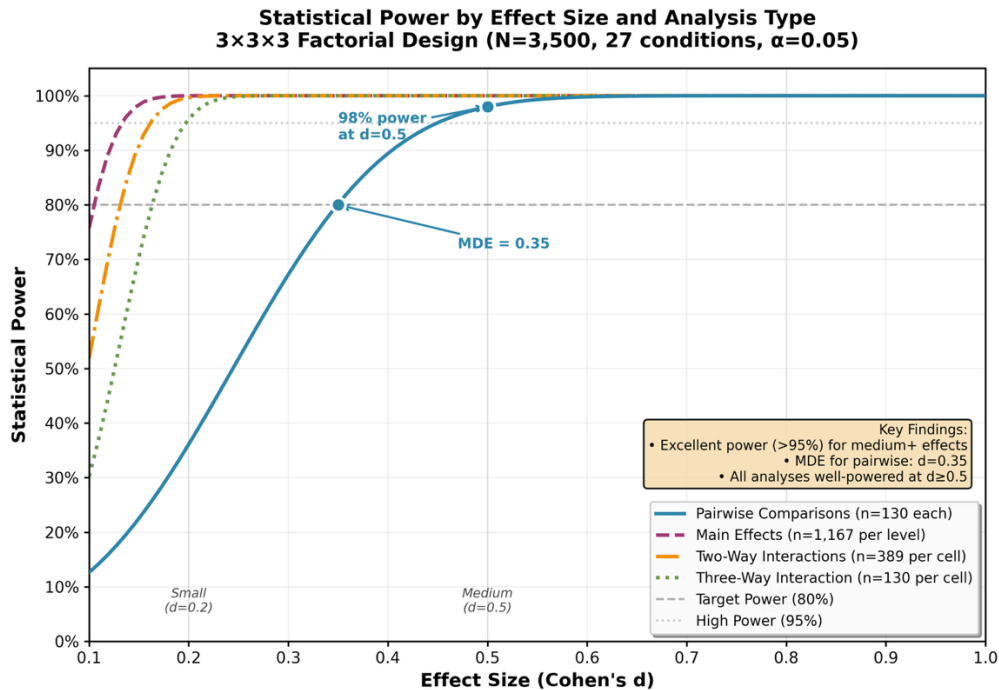
Each respondent sees 1 of 27 randomly assigned vignette conditions then answer 6 outcome questions (C18-23): C18: Democracy satisfaction (4-point scale), C19: trust in investigation (4-point scale), C20: trust in election outcome (4-point scale), C21: trust in the integrity of the election (4-point scale), C22: regulation preferences (4-point scale), C23: concerns in Australia (4-point scale).

### 5.3 Statistical Power

Figure 1 displays statistical power as a function of effect size (Cohen's  $d$ ) for different analysis types in our 3x3x3 factorial vignette experiment with  $N=3,500$  respondents distributed across 27 experimental conditions (approximately 130 respondents per condition). The power curves demonstrate that our design achieves excellent statistical power (>95%) for detecting medium or larger effects ( $d \geq 0.5$ ) across all analysis types, including pairwise comparisons between specific conditions (98% power), main effects of individual factors (>99% power), two-way interactions (>99% power), and the three-way interaction (>99% power). For pairwise comparisons—our

most conservative test with the smallest effective sample size—the minimum detectable effect with 80% power is  $d=0.35$ .

**Figure 1: Statistical Power Calculation**



## 6. ANALYSIS

### 6.1 Primary Analysis

We will employ factorial Analysis of Variance (ANOVA) to test the main effects and interactions of our three experimental factors (investigation authority, prevalence level, voter identification level) on key outcome variables including democratic satisfaction, trust in investigations, perceived AI influence on elections, and regulatory preferences. The factorial ANOVA framework will allow us to decompose total variance in outcomes into components attributable to each factor independently (main effects), combinations of two factors (two-way interactions), and the joint effect of all three factors (three-way interaction), while testing whether mean differences across experimental conditions exceed what would be expected by chance. We will report coefficients with 95% confidence intervals.

### 6.2 Effect Size Interpretation and Substantive Significance

Given that all outcome variables are measured on 4-point scales (ranging from 1 to 4), we will interpret effects using both standardized (Cohen's  $d$ ) and raw metric differences, with the full scale representing 100% of the possible range (3 points). Following conventional guidelines, we consider Cohen's  $d$  values of 0.2, 0.5, and 0.8 as small, medium, and large effects respectively. However, given the ordered categorical nature of our outcomes, we prioritize interpretation of raw differences in mean responses. We consider an effect substantively meaningful if the experimental manipulation shifts average responses by at least 0.3 points on the 4-point scale, representing approximately 10% of the total scale range. For example, an authority

manipulation that shifts trust in investigation findings from "trust somewhat" (mean = 2.5) to 2.8 would represent a 10% shift along the trust continuum, suggesting movement toward the "trust completely" category. Effects of 0.5 points (17% of scale range) or larger would be considered moderate to large substantive effects. We will also report the percentage of respondents crossing key thresholds (e.g., moving from trust to distrust categories) to provide intuitive measures of practical significance. These thresholds are specified prior to data analysis to distinguish effects that matter for democratic attitudes and regulatory preferences from those that, while statistically detectable, represent minimal attitudinal change.

### **6.3 Subgroup Analyses**

We will explore heterogeneous treatment effects by:

- Age, sex and education
- Political interest
- Political trust
- Political attitudes
- Party identification

### **6.4 Robustness Checks**

To assess the robustness of our findings, we will conduct several sensitivity analyses. First, we will estimate alternative specifications. Second, we will address missing data through multiple imputation. Third, we will conduct sensitivity analyses excluding respondents who fail attention checks or exhibit suspicious behaviour (e.g., straightlining, completion times below threshold), ensuring our results are not driven by low-quality responses.

## **7. ETHICS**

Data collection for the survey received ethical approval from the ANU Human Research Ethics Committee (approval 2021/430).

## **8. TIMELINE AND REPORTING**

Data analysis will begin immediately upon completion of data collection in December 2025. Results will be reported in accordance with the pre-registered hypotheses, with additional analyses clearly reported. Null findings will be reported. All materials, data, and code will be made available in accordance with transparency and reproducibility standards.



## 9. SURVEY QUESTIONS DECEMBER 2025

### MODULE: AI-Generated Political Information and Democracy

#### Authors:

Prof. Nicholas Biddle, Australian National University

Dr. Svitlana Chernykh, Australian National University

Dr. Constanza Sanhueza Petrarca, Australian National University

*Note: This section presents the AI module only. The survey includes additional questions that will serve as control variables, including demographic characteristics (age, gender, education), political attitudes, and measures of AI familiarity and use.*

### A. AI AND DEMOCRACY

\*(ALL)

C8. Overall, how much of a threat do you think generative AI is to democracy?

Please use a scale from 0 to 10, where 0 means "not a threat at all" and 10 means "a severe threat."

- 0. 0 – Not a threat at all
- 1. 1
- 2. 2
- 3. 3
- 4. 4
- 5. 5 - Neutral
- 6. 6
- 7. 7
- 8. 8
- 9. 9
- 10. 10 – A severe threat

98. Not sure

99. Prefer not to say

\*(ALL)

C9. In your opinion, does AI-generated content makes it harder or easier to find reliable political information?

- 1. It makes it much harder
- 2. It makes it somewhat harder
- 3. It makes it neither harder nor easier
- 4. It makes it somewhat easier
- 5. It makes it much easier

98. Not sure

99. Prefer not to say

\*(ALL)

C10 How much of the political content that you saw during the 2025 federal election campaign do you think was AI-generated?

1. Almost all of it
2. More than half
3. About half
4. Less than half
5. Almost none of it

98. Not sure
99. Prefer not to say

\*(ALL)

C11 How concerned are you about the following possible AI threats to Democracy?

(RANDOMISE STATEMENTS)

- a) AI-generated false/misleading information to manipulate political processes
- b) Flooding social media with biased AI content creating echo chambers
- c) Highly realistic AI-generated videos/images making people appear to say/do things they never did
- d) AI systems exhibiting unequal treatment based on race, gender, or political characteristics
- e) AI systems restricted by specific entities such as governments or political institutions, limiting transparency and reinforcing existing power structures
- f) AI systems that memorize and leak sensitive personal data without consent

1. Very concerned
2. Somewhat concerned
3. Not too concerned
4. Not concerned at all

98. Not sure
99. Prefer not to say

\*(ALL)

C12. In your view, how much of a risk does AI pose to the following aspects of elections?

(RANDOMISE STATEMENTS)

- a) Fair and balanced media reporting of parties and candidates
- b) Misuse of voter data
- c) Foreign interference in national elections
- d) Helping extreme actors and parties spread their message
- e) Efforts to discourage voters' turnout
- f) Efforts to manipulate voters' opinions

- g) Accuracy of information about voting procedures
- h) Fair vote counting
- i) Fair reporting of election results

- 1. A lot of risk
- 2. A moderate amount of risk
- 3. Little risk
- 4. No risk at all

- 98. Not sure
- 99. Prefer not to say

## **B. REGULATIONS**

\*(ALL)

C13 To the best of your knowledge, how would you describe the current state of generative AI regulation in Australia?

- 1. Highly regulated - Strict laws and enforcement are in place
- 2. Moderately regulated - Some regulations exist but enforcement varies
- 3. Minimally regulated - Few regulations are in place
- 4. Unregulated - No specific regulations exist

- 98. Not sure
- 99. Prefer not to say

\*(ALL)

C14 Do you think Australia should regulate AI more strictly, less strictly, or maintain current levels of regulation?

- 1. Much more strictly
- 2. Somewhat more strictly
- 3. Maintain current levels
- 4. Somewhat less strictly
- 5. Much less strictly

- 98. Not sure
- 99. Prefer not to say

\*(ALL)

C15. Next, we want to ask you questions about the use of AI for generating political content including content about voting procedures, election campaigns, political parties, candidates, policies and the government.

How much responsibility should each of the following have for regulating AI-generated political content?

(RANDOMISE STATEMENTS)

- a) The Australian government

- b) The Australian Electoral Commission
- c) Social Media Platforms
- d) Companies that develop AI
- e) News organisations and journalists
- f) Political Parties
- g) Independent fact-checking organisations
- h) International organisations
- i) The US Government
- j) The Chinese Government

- 1 A lot of responsibility
- 2 A moderate amount of responsibility
- 3 Little responsibility
- 4 No responsibility at all

98. Not sure

99. Prefer not to say

\*(ALL)

C16. The following are different ways that generative AI could be regulated during election campaigns. For each one, please indicate the extent to which you would support or oppose this approach in Australia.

(RANDOMISE STATEMENTS)

- a) Allow all political parties and candidates to use AI-generated content without restrictions
- b) Require all AI-generated political content to be clearly labelled or disclosed
- c) Allow AI use only for certain purposes but not for creating campaign content
- d) Ban AI-generated deepfakes or manipulated videos/images of candidates
- e) Prohibit all use of generative AI during the official campaign period
- f) Allow citizens and independent groups to use AI, but ban political parties from using it
- g) Analyzing voter data to target campaign messages
- h) Creating fake news articles or misleading information about opponents
- i) Creating information about voting procedures
- j) Prohibit AI-generated content during the final weeks before election day
- k) Ban all use of AI-generated content by foreign actors in Australian elections
- l) Require political parties to report their use of AI tools to the Australia Electoral Commission

- 1. Strongly support
- 2. Somewhat support
- 3. Neither support nor oppose
- 4. Somewhat oppose
- 5. Strongly oppose

98. Not sure

99. Prefer not to say

\*(ALL)

C17a. To what extent do you agree or disagree that programs should be introduced to educate the following groups in Australia about how to identify **AI-generated political content**?

(RANDOMISE STATEMENTS)

- a) Students in primary school (aged around 5 to 12 years)
- b) Students in high school (aged around 12 to 17)
- c) Students at university
- d) Adults across the age distribution
- e) Younger adults who are just starting to vote (aged around 18 to 24)
- f) Older adults who may be less familiar with Generative AI (aged around 55 years and over)
- g) Journalists and others in the media
- h) Public servants or those designing legislation

1. Strongly agree
2. Agree
3. Neither agree nor disagree
4. Disagree
5. Strongly disagree

98. Not sure

99. Prefer not to say

## C. VIGNETTE EXPERIMENT

Quota allocations for:

- <Investigating authority>:
  - the Electoral Commission – 33%
  - independent fact-checking organizations – 33%
  - a Social Media Platform – 33%
- <Prevalence Level>:
  - dominated online content related to the election. – 33%
  - was found in a moderate amount of online content related to the election. – 33%
  - was only minimally found in online content related to the election. – 33%
- <Voter identification level>:
  - Almost no voters could identify AI-generated political content – 33%
  - Some voters could identify AI-generated political content. – 33%

- Most voters could identify AI-generated political content. – 33%

[Display the following text in a grey text box, as shown here]

We want to ask you about a hypothetical situation about the use of AI in election campaigns.

*Imagine that a federal election took place in a hypothetical country. After the election, [INVESTIGATING AUTHORITY] conducted an investigation into the use of AI to generate political content during the election campaign. The investigation found that AI generated false information [PREVALENCE LEVEL]. Additionally, the investigation revealed that [VOTER IDENTIFICATION LEVEL].*

\*(ALL)

C18. Thinking about the scenario you just read, how satisfied would you be with the state of democracy in that hypothetical country?

1. Very satisfied
2. Fairly satisfied
3. Not very satisfied
4. Not at all satisfied

98. Not sure
99. Prefer not to say

\*(ALL)

C19. How much would you trust the findings of the investigation about AI-generated political content described in this scenario?

1. Trust completely
2. Trust somewhat
3. Trust a little
4. Do not trust at all

98. Not sure
99. Prefer not to say

\*(ALL)

C20. In this scenario, how likely would you think AI-generated false information influenced the election outcome?

1. Very likely
2. Somewhat likely
3. Not very likely
4. Not likely at all

- 98. Not sure
- 99. Prefer not to say

\*(ALL)

C21. How confident would you be in the integrity of the election results in this scenario?

- 1. Very confident
- 2. Somewhat confident
- 3. Not very confident
- 4. Not confident at all

- 98. Not sure
- 99. Prefer not to say

C22. Do you think that authorities in this country should regulate AI more strictly, less strictly, or maintain current levels of regulation?

- 1. Much more strictly
- 2. Somewhat more strictly
- 3. Maintain current levels
- 4. Somewhat less strictly
- 5. Much less strictly

- 98. Not sure
- 99. Prefer not to say

C23. How concerned would you be if a similar situation occurred in Australia?

- 1. Very concerned
- 2. Somewhat concerned
- 3. Not very concerned
- 4. Not concerned at all

- 98. Not sure
- 99. Prefer not to say