SUBMISSION TO THE ROYAL COMMISSION INTO ANTISEMITISM AND SOCIAL COHESION

SCHEDULE 2, *COMBATTING ANTISEMITISM, HATE AND EXTREMISM (FIREARMS AND CUSTOMS LAWS) ACT 2026*

Structural Parallels to the Robodebt Scheme and Implications for Democratic Accountability

---

Submitted to:

The Honourable Virginia Bell AC
Royal Commissioner
Royal Commission into Antisemitism and Social Cohesion

Submission Focus:

Examination of sections 92E–92G of the Australian Security Intelligence Organisation Act 1979 (ASIO Act), as enacted by Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*

Terms of Reference:

This submission addresses the Commission's mandate to examine:

- The prevalence, nature, and causes of antisemitism in Australia
- Impacts on social cohesion
- What more should be done to protect the Jewish community and promote social cohesion

Date: March 2026

Table of Contents

SUBMISSION TO THE ROYAL COMMISSION INTO ANTISEMITISM AND SOCIAL COHESION....1
SCHEDULE 2, COMBATTING ANTISEMITISM, HATE AND EXTREMISM (FIREARMS AND CUSTOMS LAWS) ACT 2026....1
Structural Parallels to the Robodebt Scheme and Implications for Democratic Accountability 1
EXECUTIVE SUMMARY....11
Relevance to the Commission's Terms of Reference....11
How This Submission Addresses the Terms of Reference....11
Why Institutional Design Matters to Antisemitism and Social Cohesion....12
Overview....12
I. The Structural Parallel to Robodebt....13
Five Fatal Features—Then and Now....13
The Critical Distinction....13
II. What Parliament Was Not Told....13
The Robodebt Omission in the Explanatory Memorandum....13
What Parliamentary Debate Revealed....14
III. Why This Matters: The Political Sustainability Argument....14
Tools Available to Future Governments....14
The 2026 Census and Population Profiling Risk....14
IV. Critical Provisions Requiring Examination....15
Section 92E: Algorithmic Assessment Authority....15
Section 92F: Ministerial Scope Control....15
Section 92G(3): Immunity from Invalidity....15
V. Why Suspend Now, Not Later....16
The Entrenchment Problem....16
Prevention vs Remediation....16
VI. Core Recommendations....16
Tier 1: Urgent Action Required....16
Tier 2: Structural Reform....17
VII. What This Submission Does Not Argue....17
What This Submission Does Argue....17
VIII. Commission's Opportunity....18
Conclusion....18
OPENING NOTE TO THE ROYAL COMMISSION....19
Acknowledgement of Tragedy....19
Personal Statement: A Lifelong Commitment....19
Context: What This Submission Anticipates....20
What This Submission Argues For....21
A Note on Tone....22
NOTE: THIS IS SUBMISSION 1 OF 3....23
Submission 1 (This Document): Systemic Patterns and Institutional Restraint....23
Submission 2 (Planned): Constitutional Reform for Social Cohesion....23
Submission 3 (Planned): Israel/Palestine - A European Union Model for Peace....23
Why Three Submissions?....24
Independence of Submissions....24
Timeline Note....24
EXECUTIVE Overview....26

TO THE ROYAL COMMISSION INTO ANTISEMITISM AND SOCIAL COHESION.....26
THE CORE FINDING THIS SUBMISSION URGES.....26
THE CONCEALMENT STRATEGY.....27
What the EM Concealed.....27
The Pattern of Concealment.....28
THE PARLIAMENTARY FAILURE.....29
Hansard Analysis.....29
What This Demonstrates.....29
THE SYSTEMIC PATTERN: TAMPA TO SCHEDULE 2 (2001-2026).....30
The Trajectory.....30
Hannah Arendt's Warning as an Interpretive Lens.....32
THE CHOICE BEFORE THE COMMISSION.....33
Option A: Post-implementation review following potential harm.....33
Option B: Intervene Now, Prevent Repetition.....33
The Political Sustainability Question.....33
Tools That Outlast Their Creators.....34
Not an Unintended Consequence.....34
Why This Submission Is Relevant to the Commission.....35
MAPPING TO ROYAL COMMISSION TERMS OF REFERENCE.....36
Scope of Contribution.....36
PART I: WHAT IS ANTISEMITISM? THE DEFINITIONAL CONTEST.....37
Introduction: The Problem of Contested Definitions.....37
The Contemporary Definitional Landscape.....37
1. IHRA Working Definition (2016).....37
2. Jerusalem Declaration on Antisemitism (JDA) (2021).....38
3. Nexus Document (2021).....38
4. Project Esther (2024).....38
Areas of Consensus and Divergence.....39
The Institutional Challenge: When Definitions Compete.....39
The "Reasonable Member of the Target Group" Standard.....39
Evidence of Divergent Interpretations in Australian Public Debate.....40
Implications for the Antisemitism Bill.....40
Rule-of-Law Considerations.....41
Conclusion.....41
PART II.....42
Structural Design, Escalation Risk, and Institutional Safeguards.....42
Overview: From Social Harm to Institutional Architecture.....42
I. Violence, Attribution, and Legislative Reaction.....42
A. Complexity of Public Violence.....42
B. The Bondi Incident: Limits of Present Knowledge.....44
C. The Risk of Reactive Institutional Expansion.....44
II. Protective Measures and the Risk of Structural Escalation.....44
III. Administrative Classification and the "Right to Have Rights".....45
IV. Automation, Contestability, and Administrative Insulation.....46
V. Definitions, Moral Language, and Institutional Design.....46
VI. The Security–Liberty Trade-Off.....47
VII. Political Volatility and Institutional Durability.....47
Part II Conclusion.....48
Part III: Antisemitism Is Real.....49
Introduction: Acknowledging the Reality of Harm.....49

Contemporary Context....49
Far-Right Extremism....49
Recorded Incidents....49
Public Demonstrations Since October 2023....50
Large-Scale Gaza Solidarity Demonstrations....50
Counter-Demonstrations and Far-Right Presence....50
Instrumentalisation and Narrative Competition....50
The Risk of Misdirected Policy Responses....51
The Importance of Distinguishing Jews from State Policy....51
Why Careful Framing Matters....52
Conclusion: Affirming Reality While Questioning Method....52
PART IV: LANGUAGE, MORAL FRAMING, AND INSTITUTIONAL POWER....53
Introduction: Why Language Matters in Moments of Fear....53
The IHRA Preamble and the Installation of a Moral Frame....53
What Moral Framing Does Inside Institutions....53
1. From Conduct to Identity....54
2. The Elastic Quality of Moral Categories....54
3. Perceived Existential Threat and Administrative Expansion....54
A Thought Experiment: Substituting Moral Language with Legal Language....55
Language and Institutional Behaviour: Historical Patterns....55
Why Universities and Public Institutions Become Particularly Vulnerable....55
The Liberal-Democratic Alternative: Legal Language and Institutional Restraint....56
Narrative Framing, Motive Attribution, and Institutional Acceleration....56
Connection to the Broader Submission....58
Transitional Bridge: From Language to Institutional Design....58
Part V: The Racial Framework Problem....59
Introduction: When Protection Relies on Fixed Identity....59
Identity as Fluid Rather Than Fixed....59
Racism as a Colonial Concept....60
The Race–Religion Distinction in Evidence to the Committee....60
Descent-Based Identity Frameworks Across Communities....61
The Ethnocultural Categorisation in Law....61
The Practical Legal Consequence: Differential Protection....62
The Problem of Legal Essentialism....63
Identity Diversity Across Communities....63
Human Rights Without Racial Hierarchy....64
The Risk of Re-creating Colonial Logic....64
Conclusion: Protection Without Essentialism....64
PART VI: INSTITUTIONAL DESIGN AND HISTORICAL LESSONS....66
Introduction: Why Machinery Matters More Than Motive....66
Hannah Arendt: Historical Context and Analytical Framework....66
Background and Scholarly Contribution....66
Core Insights from Arendt's Analysis....67
1. Persecution at Scale Requires Administrative Infrastructure....67
2. The "Banality of Evil" and Bureaucratic Diffusion....68
3. Rightlessness and Legal Limbo....69
4. The Automation of Administrative Processes....69
Application to Contemporary Policy Design....70
Not Prediction—Pattern Recognition....70
Lessons from Robodebt....71

The Royal Commission's Findings....71
Schedule 2 Through an Institutional Design Lens....72
Structural Similarities to Robodebt....72
The Question of Protective Intent....72
Why Good Intentions Are Insufficient....72
Why This Matters for Social Cohesion....73
Trust in Institutions Depends on Perceived Fairness....73
Conclusion: Learning from Historical Patterns....74
PART VII: HISTORICAL PATTERNS IN AUSTRALIAN ADMINISTRATIVE POLICY....76
Introduction: Examining Recurring Structural Features....76
I. The Tampa Incident and the Pacific Solution (2001)....76
Historical Context....76
Political Dynamics and Narrative Construction....77
The Political Environment: One Nation and Electoral Strategy....77
Bipartisan Convergence....77
Administrative Structure: The Pacific Solution Framework....78
II. Offshore Detention: Persistence Across Governments (2001–2025)....78
Continuity Despite Political Change....78
Financial and Human Costs....78
Corruption and Procurement Concerns....79
International Legal and Diplomatic Dimensions....80
The Question of Ideological Commitment....80
III. Robodebt: Automated Decision-Making Without Safeguards (2015–2019)....81
The Income Averaging Mechanism....81
Procedural Concerns....81
Documented Harms....81
Accountability Outcomes....81
Recommendation 17.1: A Warning Against Automation Without Safeguards....82
IV. Schedule 2: Replication 28 Months Later (2023–2026)....82
The Legislative Framework....82
Similarities to Robodebt Architecture....82
Parliamentary Awareness and Concealment....83
Inspector-General of Intelligence and Security Warning....83
V. Pattern Analysis: Recurring Structural Features....84
Common Design Elements Across Cases....84
The Question of Causation....84
Arendt's Framework: Techniques and Rightlessness....85
VI. Social Cohesion Implications....85
Trust and Institutional Legitimacy....85
The Deterrence Function and Community Division....86
VII. Why Royal Commission Into Offshore Detention May Be Necessary....86
The Accountability Deficit....86
What a Royal Commission Could Examine....86
Acknowledging Political Unlikelihood....87
VIII. Connection to Schedule 2 and Current Legislative Response....87
The Pattern's Relevance to Antisemitism Policy....87
Breaking the Pattern....88
Conclusion: The Commission's Opportunity....88
PART VIII: SCALE AND INSTITUTIONAL RISK....90
The Fundamental Question for Liberal Democracy....90

Individual Violence and Institutional Power....90
Scale and the Asymmetry of Power....90
Lessons from Recent Australian Experience....91
Evaluating Legislative Risk in Context....92
Liberal Democratic Risk Management....92
Why Institutional Design Matters....93
Institutional Risk and Social Cohesion....93
The Commission’s Role....93
Conclusion: A Liberal-Democratic Principle....94
PART IX: PARLIAMENTARY CONSIDERATION AND LEGISLATIVE TRANSPARENCY....95
Introduction: Examining the Information Available to Parliament....95
I. The Robodebt Context and Selective Application in the Explanatory Memorandum....95
The Robodebt Royal Commission Findings (July 2023)....95
The Asymmetry in the Explanatory Memorandum....96
The Unexplained Differential Treatment....96
Questions Raised by the Asymmetry....97
What This Asymmetry Suggests....97
II. Characterisation of the Automation Framework in the Explanatory Memorandum....98
How the Explanatory Memorandum Describes Schedule 2....98
Features Not Prominently Disclosed....98
The "Firearms Licensing" Frame....99
III. Parliamentary Debate: Analysis of Hansard Records....99
Overview of Debate....99
Primary Focus of Debate....100
Statistical Analysis of Hansard Content....100
The Single Exception: Senator McDonald's Statement....101
What Members Appear to Have Understood They Were Enacting....101
IV. The Parliamentary Committee Process....102
Timeline and Constraints....102
UN Special Rapporteur's Concern....103
Committee Report Findings....103
V. Last-Minute Constitutional Architecture Change: State Law Incorporation....103
The Bill Split and Substantive Amendment....103
What Changed: Federal Incorporation of State Racial Hatred Laws....103
Professor Anne Twomey's Constitutional Analysis....104
Constitutional Questions Raised....104
Timing and Democratic Process Concerns....106
VI. The Inspector-General's Warning....106
IGIS Submission to Parliamentary Committee....106
Parliamentary Response....106
VII. International Comparisons: Transparency Standards for Automated Government Systems 107
European Union AI Act....107
United Kingdom Algorithmic Transparency Standard....107
Schedule 2 in Comparative Context....107
VIII. Implications for Democratic Accountability....108
The Gap Between Statutory Effect and Parliamentary Understanding....108
Standards for Legislative Transparency....108
Comparative Assessment....108
IX. Connection to Social Cohesion....109

How Legislative Process Affects Community Trust....109
Specific Concerns for Affected Communities....109
Conclusion: Questions for the Commission's Consideration....110
PART X: SCHEDULE 2 STATUTORY FRAMEWORK—DETAILED ANALYSIS....112
Introduction: Examining the Automated Assessment Architecture....112
I. Section 92E: Authorising Automated Assessment Processes....112
The Core Delegation Provision....112
Legal Effect of Automated Decisions....113
Retrospective Substitution Authority....113
II. Section 92F: Ministerial Definition of Scope....114
Non-Legislative Instrument Control....114
III. Section 92G: Safeguards and Immunity Provisions....115
The "Reasonable Steps" Requirement....115
The Statutory Note....115
The Immunity Provision....116
IV. What the Statutory Framework Does Not Include....117
Absence of Specified Transparency Requirements....117
Comparison to Robodebt Royal Commission Recommendation 17.1....117
V. Comparative Analysis: International Approaches....118
European Union AI Act....118
United Kingdom Algorithmic Transparency Standard....118
Schedule 2 in Comparative Context....118
VI. Potential Applications Beyond Initial Context....119
The Statutory Language Permits Broad Application....119
Mission Creep and Scope Expansion....119
VII. Application to the Commission's Mandate....120
Relationship to Social Cohesion....120
Relationship to Antisemitism Response....120
Conclusion: Questions for Systemic Reform....120
PART XI: PRIVATE SECTOR INVOLVEMENT IN AUTOMATED ASSESSMENT....122
Introduction: Examining Potential Third-Party Deployment....122
I. The Statutory Language Regarding Third-Party Systems....122
"Arrange for the Use" — Textual Analysis....122
Oversight Requirement....123
II. Recent Developments: Security Clearances for Technology Providers....123
Palantir Technologies Security Clearance (November 2025)....123
What Protected-Level Clearance Permits....124
III. Palantir's Operations in Other Jurisdictions....124
International Context....124
United States: Immigration Enforcement Systems....124
International Development Contexts....125
Corporate Statements and Public Positioning....125
IV. Existing Australian Government Relationships....125
Contract History....125
Lobbying Activities....126
V. Statutory Framework and Private Sector Deployment....126
What Schedule 2 Permits Regarding Third-Party Systems....126
What Schedule 2 Does Not Require Regarding Third-Party Systems....127
Comparison to Public Sector Development....127
VI. International Comparisons: Regulation of Private Sector AI in Government....128

European Union Approach....128
United Kingdom Framework....128
Schedule 2 in Comparative Context....128
VII. Questions for Democratic Accountability....128
Transparency Concerns....128
Oversight Challenges....129
Parliamentary Control....129
VIII. Alternative Approaches: Safeguards for Third-Party Deployment....130
If Private Sector Involvement Proceeds....130
International Best Practice Models....130
IX. Connection to Social Cohesion and Commission Mandate....131
Public Trust in Automated Systems....131
Specific Concerns for Protected Communities....131
Conclusion: Questions Requiring Consideration....131
Part XII Additional AI-Specific Risks in Intelligence Assessment Contexts....133
Operational Consequences of Analytical Errors....133
The "Hallucination" Problem in AI-Assisted Analysis....134
Professional Incentives in Classified Environments....134
Magnitude and Reversibility Concerns....135
The Efficiency Paradox: When Correction Costs Exceed Automation Benefits....135
The Robodebt Precedent on Correction Costs....136
The Irrecoverable Case....137
Design Implication....137
Design Implication....138
The Oversight Challenge: When Performance Reporting Is Also Automated....138
Design Implication for Oversight....140
The 2026 Census: A High-Risk Data Matching Event....141
Connection to Recommendation 17.1....144
PART XIII: RECOMMENDATIONS TO THE ROYAL COMMISSION....145
Overview: A Tiered Framework....145
Introduction: From Analysis to Action....145
RECOMMENDATION 1 (TIER 1): PRIMARY — REPEAL OR SUBSTANTIALLY AMEND SCHEDULE 2 CHANGES TO THE ASIO ACT (SECTIONS 92E–92G)....146
A. The Finding This Submission Urges....146
B. The Recommendation This Submission Urges....146
C. The Political Sustainability Argument....147
D. Interim Measure (Non-operationalisation)....148
RECOMMENDATION 2 (TIER 1): REFERRALS FOR INVESTIGATION AND REVIEW..148
A. NACC....148
B. IGIS....148
C. CDPP / AFP....148
D. APSC....148
E. Solicitor-General Advice....148
RECOMMENDATION 3 (TIER 2): MINIMUM LEGISLATIVE SAFEGUARDS IF ANY AUTOMATION FRAMEWORK REMAINS....149
RECOMMENDATION 4 (TIER 2): PARLIAMENTARY PROCESS AND EXPLANATORY MEMORANDUM REFORMS....149
RECOMMENDATION 5 (TIER 2): ROBODEBT RECOMMENDATION 17.1 COMPLIANCE MONITORING....150

RECOMMENDATION 6 (TIER 2): ESTABLISH ALGORITHMIC ACCOUNTABILITY OVERSIGHT....150
RECOMMENDATION 7 (TIER 1): PROTECT 2026 CENSUS DATA FROM ALGORITHMIC PROFILING....150
Immediate Safeguards Required:....151
Rationale:....151
Connection to Jewish Community Protection:....151
RECOMMENDATION 8 (TIER 2): COMMUNITY CONSULTATION AND PUBLIC EDUCATION....152
RECOMMENDATION 9 (TIER 2): INTERNATIONAL BENCHMARKING....152
RECOMMENDATION 10 (TIER 2): IMPLEMENTATION PATHWAY....152
Conclusion....152
PART XIV — Alternative Framework for Accountable Algorithmic Governance....154
Introduction: Demonstrating the Possible....154
I. CORE PRINCIPLES FOR LEGITIMATE ALGORITHMIC GOVERNANCE....155
Principle 1: Human Decision-Making Preserved....155
Principle 2: Notification Before Adverse Action....155
Principle 3: Methodology Transparency, Source Protection....156
Principle 4: Independent Review on the Merits....156
Principle 5: No Immunity from Invalidity....157
Principle 6: Parliamentary Definition of Scope....157
Principle 7: Prohibition on Retrospective Operation....157
Principle 8: Government Development and Operation Only....158
II. Illustrative Legislative Model (Non-Prescriptive)....158
A. Proposed Legislative Structure....158
B. Key Provisions (Illustrative)....158
Part 1: Preliminary....158
Part 2: Algorithmic Assessment Framework....159
Part 3: Transparency and Accountability....160
Part 4: Independent Review....161
Part 5: Safeguards and Limitations....162
Part 6: Parliamentary Oversight....163
III. TRANSPARENCY ARCHITECTURE: BALANCING SECURITY AND ACCOUNTABILITY....164
A. The False Dichotomy....164
B. International Models Demonstrating Feasibility....164
C. Australian Precedent - Existing Automated Systems with Transparency....165
D. Proposed Transparency Levels for Intelligence Assessment....166
IV. ACCOUNTABILITY MECHANISMS: MAKING OVERSIGHT REAL....167
A. Multi-Layered Oversight Model....167
B. Special Advocate Model for Classified Material....168
C. Bias Testing and Remediation Requirements....169
V. IMPLEMENTATION PATHWAY: FROM HERE TO THERE....170
Phase 1 - Immediate (0-6 months): Stop the Harm....170
Phase 2 - Short Term (6-12 months): Develop Replacement Framework....170
Phase 3 - Medium Term (12-24 months): Build Capability....171
Phase 4 - Long Term (2-5 years): Full Operation with Oversight....171
VI. INTERNATIONAL COMPARISON: LEARNING FROM OTHERS....172
A. European Union - AI Act (2024)....172
B. United Kingdom - Algorithmic Transparency Standard....173

C. Canada - Algorithmic Impact Assessment....173
D. New Zealand - Algorithm Charter....174
E. What Australia Can Learn....174
VII. TECHNICAL FEASIBILITY: ADDRESSING COMMON OBJECTIONS....175
Objection 1: "Transparency enables evasion"....175
Objection 2: "Courts can't review classified intelligence"....176
Objection 3: "Human review slows down process"....176
Objection 4: "Government can't afford extensive oversight"....176
Objection 5: "Private sector more capable than government"....177
Objection 6: "Anti-discrimination testing is impractical"....177
VIII. CONCLUSION: TECHNOLOGY CAN SERVE JUSTICE....178
APPENDICES....180
APPENDIX A: TIMELINE OF KEY EVENTS....180
The Robodebt Era (2015-2019)....180
Robodebt Ending and Royal Commission (2020-2023)....181
Schedule 2 Development and Passage (2025-2026)....181
The Broader Pattern Context....183
APPENDIX B: OFFSHORE DETENTION STATISTICS....184
Population Statistics....184
Duration Statistics....184
Deaths and Serious Harm....185
Financial Statistics....185
Contractor Issues....185
International Law Considerations....186
APPENDIX C: GLOSSARY OF TERMS....186
APPENDIX D: REFERENCES AND SOURCES....188
Primary Legal Sources....188
Government and Official Sources....189
Academic and Expert Sources....189
Media and Investigative Sources....189
International Comparisons....190
NGO and Civil Society Sources....190
CONCLUSION....192
From Analysis to Action....192
The Core Argument....192
What This Submission Does Not Claim....193
The Commission's Opportunity....193
What Success Would Look Like....194
The Broader Stakes: Social Cohesion....194
A Personal Reflection....195
The Political Lifecycle of Security Powers: A Critical Warning....196
The Institutional Design Question....196
Why Political Confidence Cannot Replace Institutional Safeguards....196
The Historical Pattern....196
The Stakes for the Jewish Community....197
The Commission's Unique Responsibility....197
Final Words to the Commission....198

EXECUTIVE SUMMARY

Submission to the Royal Commission into Antisemitism and Social Cohesion

Re: Schedule 2, *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* — Structural Parallels to the Robodebt Scheme

---

Relevance to the Commission's Terms of Reference

The Royal Commission was established to examine:

1. **The prevalence, nature, and causes of antisemitism in Australia**
2. **Impacts on social cohesion**
3. **What more should be done to protect the Jewish community and promote social cohesion**

How This Submission Addresses the Terms of Reference

1. Protecting the Jewish Community (ToR 3)

This submission examines Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*—legislation enacted explicitly to protect Jewish Australians after the Bondi terrorist attack.

Core concern: Well-intentioned protective measures can create institutional vulnerabilities if their structural design replicates features of previously failed systems.

Specific relevance: The submission demonstrates that Schedule 2's automation framework (sections 92E–92G) replicates structural features condemned by the Robodebt Royal Commission 28 months earlier. **Tools created today to protect Jewish Australians must not become, in future decades, tools that could harm them.**

2. Social Cohesion (ToR 2)

Social cohesion depends on communities' trust that:

- Government powers are exercised lawfully and transparently
- Legislative processes provide adequate scrutiny and accountability
- Protective frameworks apply consistent safeguards across all communities

Submission's relevance:

- Examines whether Schedule 2's passage met transparency standards appropriate for frameworks affecting individual rights
- Analyses differential treatment: Robodebt safeguards explicitly applied to ACIC automation but omitted from ASIO automation in the same legislation
- Considers how legislative opacity and immunity provisions affect community trust in democratic institutions

3. Responses and Remedies (ToR 3 — "what more should be done")

The submission provides:

- **Detailed analysis** of Schedule 2's structural design and operational implications
- **Specific recommendations** for suspension, review, and redesign before operational entrenchment
- **Alternative framework** (Part XIV/Appendix) demonstrating how automation could be implemented lawfully with Robodebt Royal Commission safeguards
- **Preventive approach**: addressing structural vulnerabilities before harm occurs, not after

Why Institutional Design Matters to Antisemitism and Social Cohesion

The submission argues that **the Jewish community's long-term safety depends on institutional constraints that bind all future governments**, not only on current government intentions.

Historical lesson: Tools created in response to one crisis persist across decades and multiple governments. Australia's offshore detention regime (established 2001, dismantled 2008, re-established 2012, continuing to present) demonstrates how frameworks designed for one purpose can endure across changes in government and shifts in public opinion.

Current concern: Schedule 2 creates algorithmic assessment infrastructure that will be available to future governments of unknown composition, facing unknown threat environments, with potentially different views on who constitutes a "security risk."

Commission's opportunity: Recommend structural reforms that protect all communities—including the Jewish community—by ensuring crisis-response frameworks do not create lasting institutional vulnerabilities.

Overview

This submission examines Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*, which introduced sections 92E–92G into the Australian Security Intelligence Organisation Act 1979 (ASIO Act). These provisions authorise algorithmic intelligence assessment processes with statutory immunity from invalidity.

Core submission: Schedule 2 replicates the structural design features that enabled the Robodebt scheme's unlawful operation—28 months after the Robodebt Royal Commission condemned that architecture as fundamentally incompatible with the rule of law.

Recommendation: The Commission should recommend Parliament suspend implementation of sections 92E–92G pending comprehensive review, before these provisions become operationally entrenched and politically irreversible.

I. The Structural Parallel to Robodebt

Five Fatal Features—Then and Now

The Robodebt Royal Commission identified structural features that enabled widespread unlawful outcomes despite officials' belief they were implementing lawful policy:

Feature	Robodebt Scheme	Schedule 2 (Sections 92E–92G)
1. Automated harm without human verification	Income averaging algorithm	Algorithmic assessment determining security outcomes (s 92E(1))
2. No notification to affected persons	Debt raised without disclosure of methodology	Notification requirement deleted (s 36(1)(d) repealed)
3. No meaningful review pathway	Debt assumed valid; onus on citizen to disprove	Immunity from invalidity for safeguard non-compliance (s 92G(3))
4. Operational opacity	Methodology concealed as "commercial in confidence"	Classified intelligence + proprietary systems
5. Immunity claims	Initially claimed (later rejected by courts)	Codified in statute (s 92G(3))

The Critical Distinction

Robodebt's immunity was claimed but ultimately rejected by courts as unlawful.

Schedule 2's immunity is enacted in statute (s 92G(3)), potentially insulating automated assessment processes from the judicial review that exposed Robodebt's unlawfulness.

II. What Parliament Was Not Told

The Robodebt Omission in the Explanatory Memorandum

The Explanatory Memorandum (EM) to the Bill demonstrates government awareness of Robodebt Royal Commission findings—but applies them **selectively**:

For ACIC automation (criminal intelligence assessments, EM paragraph 287):

> "The clauses in this Division are based on the Government's response to recommendation 17.1 of the Royal Commission into the Robodebt Scheme Report and the Commonwealth Ombudsman's Automated Decision-making Better Practice Guide."

For ASIO automation (sections 92E–92G, EM paragraphs 354–362):

- ✗ No reference to Robodebt Royal Commission
- ✗ No reference to Recommendation 17.1
- ✗ No reference to Ombudsman's Automated Decision-making Guide
- ✗ No statement about compliance with administrative law principles

Question: If Robodebt safeguards are necessary for criminal intelligence automation, why not for security intelligence automation with potentially greater individual consequences?

What Parliamentary Debate Revealed

Statistical analysis of 159 pages of Hansard (combined Senate and House):

Term	Mentions	Context
Firearms/guns	400+	Buyback, licensing, regulation
Antisemitism/Jewish	300+	Attack, victims, community safety
Islamic/extremism	200+	Ideology, terrorism, preachers
Robodebt	**1**	Senator McDonald only (Senate Hansard p. 27)
AI/automation	**1**	Senator McDonald only
Algorithm	**0**	—
Section 92E/92F/92G	**0**	—
Immunity from invalidity	**0**	—

One senator identified the Robodebt connection. It generated no further parliamentary discussion.

III. Why This Matters: The Political Sustainability Argument

Tools Available to Future Governments

The Jewish community's long-term safety depends not only on current government intentions but on **institutional constraints** that bind future governments regardless of their ideology.

Historical precedent: Australia's offshore detention regime (established 2001, dismantled 2008, re-established 2012) demonstrates how tools created for one purpose persist across decades and multiple governments with varying ideological commitments. Created by a Coalition government in response to the "Tampa crisis," briefly dismantled by Labor, then re-established by Labor and expanded by subsequent Coalition governments, the regime shows how crisis-response frameworks become entrenched regardless of which party holds power.

Schedule 2 creates tools that will be available to:

- Future governments of different political composition
- Governments facing different threat environments
- Governments with different definitions of "security risk"

The 2026 Census and Population Profiling Risk

Timing concern: The August 2026 Census occurs while sections 92E–92G are newly operational.

Capability created:

- Census question 23: "What is the person's ancestry?" (up to two responses)
- 2021 Census: 99,956 Australians identified "Jewish" ancestry
- Sections 92E–92G + ministerial instrument (s 92F) = potential infrastructure for population-wide algorithmic assessment
- August 2026 timeframe = limited opportunity for democratic course correction

Question: Should tools capable of population-wide profiling based on ancestry data exist, even if current government has no intention to use them that way?

IV. Critical Provisions Requiring Examination

Section 92E: Algorithmic Assessment Authority

Authorises: Director-General may arrange for computer programs to undertake "specified assessment action"

Scope: "Specified assessment action that may, or must, be taken by the Organisation under this Act" (s 92E(1))

Analysis: Statutory language is not limited to firearms licensing. The ASIO Act authorises intelligence assessments for:

- Visa decisions (Migration Act s 36(1C))
- Citizenship determinations (Australian Citizenship Act s 17(6))
- Security clearances (multiple Acts)
- Passport decisions (Australian Passports Act s 14)

EM characterisation: "Firearms licensing efficiency"

Statutory authorisation: Potentially much broader via ministerial instrument

Section 92F: Ministerial Scope Control

Authorises: Minister may determine what constitutes "specified assessment action" by legislative instrument

Critical feature: Scope expansion does not require return to Parliament

Safeguard: Disallowance by either House—but disallowance motions historically have low success rates for technical instruments

Section 92G(3): Immunity from Invalidity

Provides: Automated assessment not invalid merely because "reasonable steps" safeguards not complied with

Effect: Creates two-tier validity regime—procedural compliance encouraged but not legally required for validity

Robodebt parallel: Royal Commission found "reasonable steps" requirements proved difficult to enforce; immunity clause removes even that incentive

---

V. Why Suspend Now, Not Later

The Entrenchment Problem

Once operational systems are deployed:

1. **Sunk costs** create resistance to change ("We've already invested $X million")
2. **Operational dependency** develops ("The system is processing 10,000 assessments monthly; we can't turn it off")
3. **Political reversal costs** increase ("Dismantling would create intelligence gaps")
4. **Bureaucratic normalisation** occurs ("This is just how we do assessments now")

Historical example: The Australian Taxation Office's income averaging algorithm (Robodebt) operated for years before legal challenge, despite early warnings from departmental legal officers.

Prevention vs Remediation

Prevention strategy (recommended):

- Suspend sections 92E–92G before operational deployment
- Conduct comprehensive review of structural design
- Implement Robodebt Royal Commission safeguards explicitly
- Return to Parliament with revised framework

Remediation strategy (alternative):

- Wait for operational deployment
- Rely on Inspector-General oversight detecting problems
- Hope affected individuals can access judicial review despite immunity clause
- Potentially face another Royal Commission in 5–10 years examining how automated systems caused harm

Robodebt lesson: Remediation after harm is devastating for victims and costly for government. Prevention is preferable.

VI. Core Recommendations

Tier 1: Urgent Action Required

Recommendation 1: Parliament should **suspend implementation** of sections 92E–92G of the ASIO Act pending comprehensive review and redesign.

Recommendation 2: Refer to appropriate bodies for investigation:

- Inspector-General of Intelligence and Security: compliance with ASIO Act safeguards
- Commonwealth Ombudsman: consistency with automated decision-making standards
- Australian Information Commissioner: Census data protection implications
- Parliamentary Joint Committee on Intelligence and Security: ongoing monitoring

Recommendation 7: Immediate measures to protect 2026 Census data from algorithmic profiling based on ancestry responses.

Tier 2: Structural Reform

Recommendation 3: If automation framework retained, minimum legislative safeguards must include:

- Explicit compliance with Robodebt RC Recommendation 17.1
- Mandatory human verification for adverse outcomes
- Restoration of notification rights
- Removal of immunity provisions
- Statutory limitation to genuinely high-volume, low-risk assessments

Recommendation 4: Parliamentary process reforms:

- Plain-language summaries mandatory for complex technical provisions
- Minimum 90-day consultation for frameworks affecting individual rights
- Explanatory Memoranda must disclose relevant Royal Commission findings
- Substantive constitutional changes require separate inquiry

---

VII. What This Submission Does Not Argue

This submission **does not** argue:

- ✗ That current government intends misuse
- ✗ That ASIO officials are acting in bad faith
- ✗ That intelligence automation is always inappropriate
- ✗ That the Bondi attack response was wrong in principle
- ✗ That antisemitism doesn't require urgent action

What This Submission Does Argue

This submission **does** argue:

- ✓ Well-intentioned responses to genuine threats can create institutional vulnerabilities
- ✓ Structural design matters more than current intentions when tools persist across governments
- ✓ Robodebt's lessons should apply to all automated government decision-making, including intelligence
- ✓ Parliament deserved to know about the Robodebt parallel before voting
- ✓ Prevention is preferable to remediation when vulnerable design features are identified
- ✓ The Jewish community's long-term safety depends on institutional constraints, not governmental goodwill

---

VIII. Commission's Opportunity

The Royal Commission into Antisemitism and Social Cohesion has an opportunity to prevent institutional vulnerability before it becomes operationally entrenched.

Key questions for Commission examination:

1. Does Schedule 2 replicate Robodebt's structural design features?
2. Why were Robodebt safeguards applied to ACIC but not ASIO automation?
3. Did Parliament have adequate information about structural parallels?
4. What prevents sections 92E–92G from scope expansion beyond firearms?
5. How does statutory immunity from invalidity (s 92G(3)) differ from Robodebt's rejected immunity claims?
6. What safeguards protect 2026 Census ancestry data from algorithmic profiling?

Broader significance:

The Commission's Terms of Reference ask: "What more should be done to protect the Jewish community and promote social cohesion?"

This submission's answer: Protect all communities—including the Jewish community—by ensuring government powers created in crisis do not persist as institutional vulnerabilities across decades.

Tools created today in the name of protecting Jewish Australians must not become, in future decades, tools that could harm them.

That is the Robodebt lesson. That is why suspension and review matter now, not later.

Conclusion

Twenty-eight months after the Robodebt Royal Commission found that automated systems with inadequate safeguards caused widespread unlawful harm, Parliament enacted Schedule 2—a framework replicating Robodebt's fatal features.

The critical question: Will Australia learn from Robodebt before harm occurs, or after?

The Commission's recommendation for suspension and comprehensive review could ensure the answer is "before."

Full submission: 20 Parts, detailed statutory analysis, international comparisons, historical patterns, and comprehensive recommendations follow in the complete document.

Contact: Robert Vose

Date: March 2026

OPENING NOTE TO THE ROYAL COMMISSION

Acknowledgement of Tragedy

This submission begins with recognition of the tragedy that prompted this Royal Commission's establishment. On the 14th of December, 2025, at Bondi Beach, sixteen people lost their lives and forty people were injured in an attack that shocked the nation and devastated their families and communities.

To the families, friends, and loved ones of those killed and injured: this submission offers deepest condolences. No analysis of policy, no examination of institutional responses, can diminish the profound loss you have suffered. Your grief is real. The fear experienced by communities is understandable. You deserve both justice for what happened and genuine protection going forward.

The Jewish community in Australia has experienced rising antisemitism—vandalism of synagogues, harassment, threats, and now this violence. These are not abstract policy questions. They are attacks on real people living their lives, practicing their faith, raising their families. This submission's critique of certain institutional responses should never be read as dismissing the reality or severity of antisemitism.

Personal Statement: A Lifelong Commitment

This submission emerges from a lifelong commitment: preventing anything like the Holocaust from happening again—to any group of people, anywhere.

The Holocaust stands as history's clearest demonstration of what happens when democracies fail — **not as a point of moral comparison, but as a warning about institutional collapse**— when persecution becomes state policy, when bureaucratic machinery is built to target a minority, when "protecting the nation" becomes justification for systematic cruelty. It must never be repeated.

Yet the lesson of the Holocaust is not merely to "protect Jews" (though we must). The lesson is "protect everyone"—because similar patterns of persecution, once established, tend to appear across different contexts and populations, and because liberal democracy's core promise is that all people, regardless of identity, deserve equal protection under law.

This is not a zero-sum game. Protecting Muslims from persecution does not threaten Jews. Protecting Jews from antisemitism does not require targeting Muslims. We do not make Jewish Australians safer by making Muslim Australians less safe. We do not protect minorities by creating new classes of marginalized people.

The rise of far-right movements globally—with their rhetoric of racial threat, their willingness to scapegoat minorities, their contempt for democratic restraints—makes protecting liberal democracy more urgent than ever. Liberal democracy, with all its frustrations and imperfections, remains the

only political system that has successfully protected religious and ethnic minorities over sustained periods. When it fails, minorities suffer first and suffer most.

This is why this submission examines not just individual acts of antisemitism, but the institutional responses—because history teaches that well-intentioned responses to genuine threats can, if poorly designed, create the very dangers they seek to prevent.

Context: What This Submission Anticipates

This Royal Commission was established amid calls for stronger laws, expanded surveillance powers, and enhanced intelligence-gathering measures. This submission proceeds on the basis that the Commission will consider proposals concerning the use of ASIO's new algorithmic assessment capabilities introduced by Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Bill 2026*, and that questions may arise as to how such capabilities could be operationalised in ways that affect particular communities, including Muslim and pro-Palestinian Australians, in the context of efforts to address antisemitism.

While Schedule 2 has been publicly justified by reference to administrative efficiency in specific contexts, including firearms licensing, the legislative framework does not, on its face, confine the algorithmic assessment capability for ASIO that it introduces to that function alone, leaving its potential application to security assessments more broadly a live question.

This submission raises concerns that such an approach risks being not merely ineffective, but dangerous.

Not dangerous because antisemitism isn't real or serious. Dangerous because:

1. **History shows that building state machinery to target minorities—even in response to genuine security concerns—creates precedents that may extend beyond original intent.** Similar structural features developed for managing "boat people" (offshore detention) later appeared in Robodebt (welfare recipients), and can now be observed in the design logic of Schedule 2, which applies across the population. These features include the placement of marginalised groups in legal limbo, the absence of notification, limited avenues for review, and systems that are insulated from accountability. Whether through direct influence, parallel bureaucratic logic, or recurring systemic patterns, the result is mission creep.
2. **Persecution of one minority does not protect another minority.** Targeting Muslim communities for enhanced surveillance, creating atmospheres of suspicion, treating political activism as a security threat—these do not make Jewish Australians safer. They create divided communities, erode social cohesion, and establish precedents that can be redirected.
3. **Liberal democracy's strength lies in restraint.** The temptation in moments of fear is to demand that government "do something"—build more powerful tools, gather more intelligence, act more decisively against threats. But unchecked power, even when initially directed at legitimate threats, becomes arbitrary power. And arbitrary power threatens everyone.
4. **Definitions of antisemitism are contested, with different frameworks offering different approaches.** The IHRA Working Definition, the Jerusalem Declaration on Antisemitism

(JDA), and the Nexus Document each reflect different community concerns and priorities. This submission examines how the IHRA's preamble language framing antisemitism as "evil" may create challenges when translated into policy—potentially creating moral rather than legal standards, where restraint becomes difficult to justify and identity considerations may overshadow conduct-based assessment. The Commission will need to consider which definitional framework best serves protective goals while maintaining liberal democratic principles, noting that the term “antisemitism” itself is contested and subject to differing interpretations.

What This Submission Argues For

This submission argues for protecting Jewish Australians through means that strengthen liberal democracy rather than undermine it:

- **Condemn and prosecute antisemitic conduct** (harassment, discrimination, threats, violence) using existing criminal and civil law—which is robust, evidence-based, and subject to review.
- **Reject frameworks that conflate political criticism with hatred**—because protecting minorities requires protecting political speech, even uncomfortable political speech. The line between criticism of Israeli government policy and antisemitism must remain clear.
- **Maintain separation between identity (protected) and ideology (criticizable)**—a person's Jewish identity is inherent and deserves absolute protection; Zionism is a political ideology and, like all ideologies, is subject to critique.
- **Refuse to build algorithmic surveillance and assessment systems with immunity from review**—not because there are no genuine security threats, but because such systems, once built, cannot be contained to their original purpose.
- **Strengthen democratic accountability rather than expand executive power**—because the best protection any minority has is a democracy with robust rule of law, independent courts, free press, and transparent government.

The question is not "should we protect Jewish Australians?" but "what means of protection serve that goal without creating new dangers?"

This submission respectfully suggests that the path forward is not more powerful, less accountable government machinery directed at particular communities, but rather:

- **Stronger rule of law** (so antisemitic conduct is consistently prosecuted)
- **Greater transparency** (so government power is visible and reviewable)
- **Robust democratic accountability** (so communities can challenge policies that harm them)
- **Social cohesion built on inclusion** (not on dividing communities into threats and protected classes)

These are not naive hopes. They are the hard-won lessons of how liberal democracies have successfully protected minorities when they've worked, and what has enabled persecution when they've failed.

A Note on Tone

This submission is long, detailed, and at times critical of government responses. This reflects the seriousness with which it takes both the problem (antisemitism and social cohesion threats) and the stakes (liberal democracy's capacity to protect minorities).

Any criticism should be read as concern for the quality and effectiveness of institutional responses, not as dismissal of the underlying problems those responses seek to address.

The goal is not to obstruct protection of Jewish Australians. The goal is to ensure that protection is achieved through means that actually work and that don't create new vulnerabilities for anyone—including, ultimately, Jewish Australians themselves, who like all minorities depend on living in a society governed by law rather than by power.

This submission proceeds from the conviction that we can protect Jewish Australians from antisemitism AND protect all Australians from arbitrary government power. These goals are not in conflict. In fact, they require each other.

NOTE: THIS IS SUBMISSION 1 OF 3

This is the first of three planned submissions to the Royal Commission into Antisemitism and Social Cohesion.

Submission 1 (This Document): Systemic Patterns and Institutional Restraint

Focus: Democratic accountability failures, the danger of "evil" language in state responses, and the importance of maintaining rule of law when protecting communities.

Core argument: Liberal democracies must respond to genuine threats (including antisemitism) through legal frameworks (criminal, unlawful, harmful) rather than moral crusade frameworks ("fighting evil"), which enable mission creep and threaten democratic foundations.

Key components:

- Analysis of Schedule 2 concealment and Robodebt architecture replication
- Documentation of 27-year pattern from Tampa → offshore detention → Robodebt → Schedule 2
- Examination of IHRA Definition "evil" framing and its political weaponisation
- 10 detailed recommendations for systemic reform
- Alternative vision showing legitimate algorithmic governance is achievable

Submission 2 (Planned): Constitutional Reform for Social Cohesion

Focus: Naturalising the Crown of Australia to strengthen national unity and address social cohesion concerns raised by the Royal Commission.

Core argument: Constitutional changes that complete Australia's sovereignty journey can address identity-based divisions and strengthen social cohesion by providing clear, unifying national framework.

Relevance to RC mandate: Directly addresses "social cohesion" aspect of Commission's terms of reference through constitutional and institutional reform proposals.

Submission 3 (Planned): Israel/Palestine - A European Union Model for Peace

Focus: Proposing an EU-like bloc structure for Israel/Palestine as long-term resolution addressing root causes of antisemitism related to Middle East conflict.

Core argument: Drawing on historical, political, and theological frameworks (with particular focus on Christian Zionist audiences in the USA), this submission will propose concrete pathways toward

sustainable peace that reduces global tensions and associated antisemitism. An approach similar to the Marshall Plan to rebuild shattered communities is proposed.

Relevance to RC mandate: Addresses international dimensions of antisemitism and provides constructive vision for resolving conflicts that fuel antisemitic rhetoric and violence.

Why Three Submissions?

The Royal Commission's mandate encompasses:

1. Antisemitism in Australia
2. Social cohesion
3. Responses and remedies

Three submissions allow comprehensive, focused responses:

Submission 1 addresses immediate concerns (institutional responses, policy frameworks) and systemic vulnerabilities (patterns enabling accountability failures)

Submission 2 addresses constitutional/structural foundations for social cohesion

Submission 3 addresses international/geopolitical dimensions and long-term peace frameworks

Together, they provide:

- **Immediate remedies** (Submission 1: how to respond to antisemitism without undermining liberal democracy)
- **Medium-term reforms** (Submission 2: constitutional changes strengthening cohesion)
- **Long-term vision** (Submission 3: international peace frameworks reducing conflict-driven antisemitism)

Independence of Submissions

While these three submissions are conceptually linked, **each stands alone**:

- Each can be evaluated independently
- Each addresses distinct aspects of the Commission's mandate
- Each provides specific, actionable recommendations
- Each draws on different evidence bases and theoretical frameworks

The Commission may:

- Accept any combination of the three
- Implement recommendations from one without others
- Use evidence from one to inform assessment of others

Timeline Note

Submission 1 (this document): Completed and ready for submission

Submission 2: In preparation, will be submitted separately

Submission 3: In preparation, will be submitted separately

Each will reference the others where relevant but will not depend on acceptance of the others.

EXECUTIVE Overview

TO THE ROYAL COMMISSION INTO ANTISEMITISM AND SOCIAL COHESION

This submission respectfully urges the Commission to investigate Schedule 2 to the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*, which inserted sections 92E-92G into the Australian Security Intelligence Organisation Act 1979 (ASIO Act).

This submission argues that Schedule 2:

- may replicate structural features identified in the Robodebt Royal Commission
- was enacted through processes that limited parliamentary scrutiny
- raises questions regarding consistency with Recommendation 17.1 of the Robodebt Royal Commission
- may reflect a longer trajectory of accountability challenges

This submission urges the Commission to recommend immediate intervention to prevent Robodebt-style harm before it occurs.

THE CORE FINDING THIS SUBMISSION URGES

Schedule 2 replicates Robodebt's unlawful architecture just 28 months after the Robodebt Commission condemned it.

Both systems share five fatal features:

Feature	Robodebt (2015-2019)	Schedule 2 (2026) amentment to ASIO Act
Automated harm	Income averaging algorithm calculated debts	Algorithmic assessment determines security risk
No notification	Recipients not told methodology	Section 36(1)(d) deleted - no notification required
No meaningful review	Debt assumed valid, burden on recipient	Section 92G(3) immunity - assessments valid even if unlawful
Opacity	Methodology concealed from recipients	Classified + proprietary algorithms - no scrutiny
Immunity claims	Government claimed debts valid (later disproven)	Statutory immunity codified in section 92G(3)

The difference: Robodebt's immunity was claimed but ultimately rejected by Federal Court. Schedule 2's immunity is codified in statute, potentially limiting avenues of judicial challenge.

The consequence: If Robodebt required a Royal Commission after 443,000 people were harmed, Schedule 2 enables harm on potentially a greater scale with even less accountability.

THE CONCEALMENT STRATEGY

This submission argues that the Explanatory Memorandum appears to have framed key provisions in ways that limited substantive engagement.

What the EM Concealed

1. Robodebt Royal Commission Findings

- Robodebt Royal Commission Report delivered 7 July 2023 (approximately 28 months before Schedule 2 was enacted)
- Recommendation 17.1 requires automated systems to comply with administrative law principles of legality, fairness, rationality and transparency
- **The Explanatory Memorandum explicitly references Recommendation 17.1 for ACIC automation** (criminal intelligence assessments for firearms licensing, paragraph 287), stating those provisions were designed to ensure compliance with administrative law principles
- **The Explanatory Memorandum makes no reference to Recommendation 17.1 for ASIO automation** (sections 92E–92G covering security assessments for visas, citizenship, clearances)
- This creates an unexplained asymmetry: Robodebt safeguards applied to one agency but not another within the same legislation
- The Inspector-General of Intelligence and Security expressly cautioned that public visibility of automated arrangements should be considered in light of the Robodebt Royal Commission, a concern not reflected in the EM's treatment of Schedule 2
  (See Inspector-General of Intelligence and Security, Submission No. 145 to the Parliamentary Joint Committee on Intelligence and Security, provided in connection with the Omnibus legislation.)

2. Immunity from Invalidity (Section 92G(3))

- Creates a statutory ‘no invalidity’ clause for failures to comply with oversight/safeguard steps (92G(1)–(2)), limiting one major pathway by which unlawful automation could be challenged.
- May violate Constitution (Chapter III separation of powers)
- **EM describes as**: "Ensures operational efficiency" (page 105 - paras 383-384)
- **EM conceals**: That this is immunity from invalidity for safeguard non-compliance

3. Private Sector Deployment (Section 92E(1) "arrange for")

- The ‘arrange for’ formulation in 92E(1) is broad enough to permit external vendor involvement in development and operation of the relevant software (subject to whatever procurement/security constraints apply).
- Palantir reportedly received Top Secret clearance in October 2025 (several months before the Bill was introduced)
- If external vendors were engaged, questions may arise regarding proprietary systems and parliamentary visibility, especially for computer program access by controversial US companies.

- Proprietary algorithms are usually exempt from parliamentary scrutiny
- National security risk (if a US corporation subject to US law is contracted for these computer programs / AI)
- **EM describes as**: "Computer programs" (page 101-103)
- **EM conceals**: Likely foreign contractor involvement and implications

4. Retrospective Substitution (Sections 92E(6)-(8))

- Allows backdating algorithmic assessments to replace previous human determinations
- Undermines vested rights and legitimate expectations
- **EM describes as**: EM presents substitution/backdating as a safeguard/correction mechanism, but it also enables substituted decisions to take effect from the date of the original automated action (92E(7)–(8)).
- **EM conceals**: That this is retrospective operation affecting settled matters

5. The 2026 Census Timing Risk

- The August 2026 Census will collect comprehensive population data including ethnicity and religious affiliation
- AI systems can match address-based Census data with other government databases (driver licenses, electoral rolls) to reconstruct individual identities
- Sections 92E–92G create legal authority for such algorithmic data matching without notification
- Once Census data is collected and potentially matched, reversal becomes extraordinarily difficult
- **EM does not address**: The timing of Census data collection relative to new algorithmic assessment powers, or safeguards to prevent Census data being used for population-wide profiling based on ancestry or religion
- **EM conceals**: The risk that Census 2026 could become the foundational dataset for automated profiling of Australian residents

The Pattern of Concealment

Every critical element enabling Schedule 2's unlawfulness was either:

- Selectively applied across agencies (Robodebt RC 17.1 for ACIC but not ASIO)
- Described in misleading euphemisms (immunity = "efficiency")
- Presented as technical or administrative updates
- Abstracted to obscure meaning (foreign contractors = "computer programs")

This submission respectfully suggests this concealment was systematic, not accidental.

The Omnibus Bill was split in a very short time-frame. The resulting split bills were not reviewed through any dedicated parliamentary inquiry. The inquiry into the Omnibus Bill was over a 48 hour window, yet the inquiry received over 7000 submissions. By any measure, this was an extraordinarily rushed exercise, and the Omnibus Bill included many diverse topics.

THE PARLIAMENTARY FAILURE

This submission documents that Parliament did not meaningfully engage with, or demonstrate understanding of, what it was enacting.

The Omnibus Bill subject to inquiry was subsequently split into two separate bills, with the split legislation and accompanying Explanatory Memoranda being made available to parliamentarians only hours before the votes were taken. The split bills introduced substantive new legislative approaches that had not been examined during the Omnibus Bill inquiry, including reliance on state and territory anti-discrimination legislation as part of the regulatory framework.

This compressed and fragmented legislative process materially limited Parliament's capacity to scrutinise the scope, operation, and consequences of the new provisions.

Hansard Analysis

(Joint sitting debate, 20 January 2026 — approximately 159 pages of Hansard)

A review of the parliamentary debate shows that discussion overwhelmingly focused on firearms and antisemitism, with little to no engagement with the substance of Schedule 2.

Topic	Mentions
Firearms	Approximately 400+ references
Antisemitism	Approximately 200+ references
Sections 92E-92G (by name)	0
"Algorithmic assessment"	0
"Immunity from invalidity"	0
Robodebt Royal Commission	1 (Senator Shoebridge)
"Automation" or "automated"	1 (Senator Shoebridge)
Palantir or private contractors	0
Notification deletion	0

Parliament's attention: Firearms and antisemitism — legitimate issues addressed elsewhere in the Bill.

Parliamentary engagement with Schedule 2: Minimal to non-existent.

The sole warning: Senator David Shoebridge (Greens) raised concerns that the Robodebt Royal Commission had recommended safeguards for automated decision-making and questioned their absence in Schedule 2.

Government response: No substantive response recorded.

Opposition response: No substantive response recorded.

Outcome: Bill passed with bipartisan support, without amendment.

What This Demonstrates

Parliament enacted Schedule 2 without:

- Engaging with the ways in which Schedule 2 departs from Robodebt Royal Commission recommendations
- Understanding the legal effect of the immunity from invalidity clause
- Recognising potential for private sector deployment authority into ASIO data systems
- Discussing algorithmic assessment scope or safeguards
- Connecting to Robodebt architecture
- Informed consent or deliberation

This submission respectfully suggests this constitutes a **massive failure of democratic administration**, echoing the language used by Commissioner Holmes in relation to Robodebt.

THE SYSTEMIC PATTERN: TAMPA TO SCHEDULE 2 (2001-2026)

This submission examines whether Schedule 2 reflects a longer trajectory of institutional responses developed over the past two to three decades.

The Trajectory

Tampa and Children Overboard (2001)

- Public debate surrounding the Tampa incident and the ‘Children Overboard’ claims contributed to a policy shift toward offshore processing and expanded executive discretion in migration control.
- The ‘Children Overboard’ claims were later found to be inconsistent with available evidence.
- The Tampa episode and the surrounding public debate were widely regarded by political commentators as significant factors in the 2001 election, after which the Howard Government expanded offshore processing arrangements.
- **Structural features that emerged**: expanded executive discretion and the placement of certain groups, such as refugees by boat, in administrative frameworks with reduced legal visibility.

Offshore Detention (2001-present, with evolving legislative frameworks)

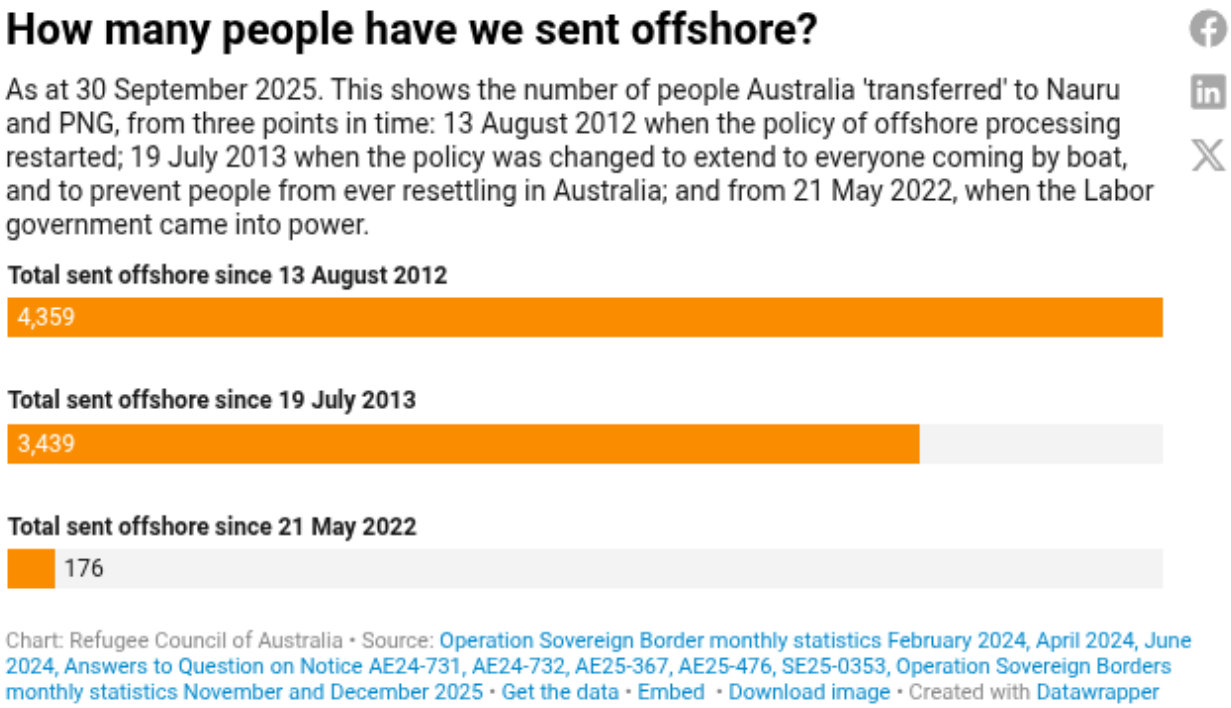

https://www.refugeecouncil.org.au/operation-sovereign-borders-offshore-detention-statistics/2/

- 4,359 people sent offshore (Refugee Council of Australia website, accessed 18/02/2026)
- 62 recorded incidents of actual self-harm (2014-2015)
  Senate Legal and Constitutional Affairs Committee
  Inquiry into Serious allegations of abuse, self-harm and neglect of asylum seekers in relation to the Nauru Regional Processing Centre, and any like allegations in relation to the Manus Regional Processing Centre, report released 21 April 2017.
- 21 deaths overall, including at least 14 deaths occurring within offshore processing environments
- $13.35 billion cost since 2012 (much more expensive per person than onshore processing) https://www.refugeecouncil.org.au/wp-content/uploads/2026/02/260213-RCOA-offshore-processing-and-arrangements-submission.pdf
- Significant controversy regarding procurement processes and contractor arrangements (including Paladin and Canstruct contracts) has been documented in parliamentary and media reporting. (Paladin: $8 company → $385M contract; Canstruct: $1.5B for 108 people)
- Findings of inconsistency with international law obligations have been raised by UN bodies, the ICC Prosecutor's office, and the PNG Supreme Court.
- **Techniques perfected**: Legal limbo, no notification, no review, immunity from normal legal process, bipartisan support preventing accountability

Robodebt (2015–2019)

- Automated income-averaging processes used to raise debts with limited explanation or investigation
- Hundreds of thousands of people affected, including approximately **381,000 individuals who later received refunds**
- Debts totalling **about $1.75 billion** written off following legal challenges
- Federal Court proceedings confirmed that averaged PAYG data could not lawfully establish debts
- Royal Commission condemned the scheme as a “massive failure” involving “venality, incompetence and cowardice”
- **Lesson demonstrated**: Automated decision-making without adequate safeguards can produce widespread administrative harm

Schedule 2 addition to ASIO Act (2026)

- Unlike the Robodebt scheme — which ultimately failed because it lacked a lawful statutory foundation — Schedule 2 embeds structural automation features that may be read as comparable to aspects of the Robodebt scheme within legislation enacted by Parliament. This shifts the issue from one of administrative unlawfulness to one of legislative design and democratic restraint. Whereas Robodebt involved the recovery of welfare debts from individuals, Schedule 2 concerns the use of algorithmic assessment within the national security framework. The significance of this difference is not merely legal but institutional: the powers created may affect individuals and communities through security assessments, depending on how they are operationalised, raising questions about safeguards, transparency, and proportionality.
- Design elements embedded in statute include: a statutory ‘no-invalidity’ framework, reduced notification requirements, and increased opacity arising from security classification.
- **This submission suggests a trajectory** in which institutional techniques first developed in migration contexts may now extend more broadly across the population.

Hannah Arendt’s Warning as an Interpretive Lens

Philosopher Hannah Arendt observed in *The Origins of Totalitarianism* (1951) that when states create administrative categories in which certain groups experience diminished legal visibility or protection, patterns of arbitrary rule may expand beyond their original targets and affect broader populations.

This submission respectfully suggests that aspects of Australia’s recent policy trajectory may be read alongside Arendt’s warning:

- **2001**: Individuals placed within offshore processing arrangements characterised by reduced access to ordinary legal protections
- **2015**: Welfare recipients subjected to automated administrative decision-making later found to lack lawful foundation (Robodebt)
- **2026**: Individuals potentially subject to algorithmic security assessment frameworks created through Schedule 2

A recurring concern identified in this analysis is that where sustained bipartisan consensus exists, traditional mechanisms of democratic accountability — including electoral rejection — may be weakened.

THE CHOICE BEFORE THE COMMISSION

This submission respectfully suggests the Commission may be presented with two broad pathways for consideration:

Option A: Post-implementation review following potential harm

- Schedule 2 implemented with the changes to the ASIO Act
- Algorithmic assessments conducted (possibly involving external contractors)
- People harmed (visas denied, employment lost, travel restricted) without notification or review
- Harm accumulates over years
- Eventually legal challenge or scandal forces accountability
- Future Royal Commission investigates harm already done
- Victims receive belated compensation and apologies

An approach analogous to the Robodebt experience — where large-scale harm was recognised only after prolonged operation and legal challenge.

Option B: Intervene Now, Prevent Repetition

- The Commission considers whether aspects of Schedule 2 changes to the ASIO Act resemble structural features identified in the Robodebt Royal Commission
- Recommends immediate non-operationalization pending review
- Parliament reconsiders with full information
- Constitutional advice obtained
- Legitimate framework developed (Part IX shows this is achievable)
- Prevents harm before it occurs
- Learns from Robodebt rather than repeating it

This is the prevention path: Use this Commission's authority to break the pattern.

The Political Sustainability Question

This submission respectfully notes a consideration that may arise during the Commission's deliberations:

The Commission may face pressure to accept sections 92E–92G as enacted, on the basis that current political alignments will ensure appropriate implementation. This submission respectfully argues that **relying on political influence rather than institutional safeguards is precisely the vulnerability these provisions create**.

Regardless of current advocacy groups' confidence in their political influence, that influence is not permanent. Institutional tools persist. Political coalitions shift.

The question before the Commission is not who controls the political narrative today, but what tools will exist when different people control government tomorrow.

Tools That Outlast Their Creators

Legislative powers do not remain in the hands of their creators. The algorithmic assessment capabilities, ministerial discretion, and immunity provisions created by sections 92E–92G will be available to future governments with different priorities.

Australia's political landscape includes far-right elements. This is documented reality: individuals displaying Nazi symbols and antisemitic messaging have demonstrated at state parliament buildings. Political parties with established records of targeting minority communities are polling at levels that could translate to increased parliamentary representation and potential coalition influence.

When political power shifts to include such elements, sections 92E–92G's capabilities will be available to them:

- Algorithmic systems that assess individuals and communities for security risk
- Ministerial discretion to expand categories without parliamentary amendment
- Immunity provisions limiting judicial challenge
- Reduced notification requirements concealing the basis of adverse determinations
- Classification frameworks preventing transparency about methodology

These features do not become inoperative when governments change. They persist.

Not an Unintended Consequence

This is not an argument about the current government's intentions or the motivations of those who drafted these provisions. This is an argument about **institutional design across political cycles**.

History demonstrates that administrative machinery created to classify and track populations based on group characteristics does not remain limited to original purposes. The machinery operates according to its structural capabilities, not the intentions of its creators.

This pattern is not an unintended consequence. It is the foreseeable operation of the tools being created.

This submission therefore respectfully argues that the Jewish community's long-term safety—and the safety of all minorities in a pluralist democracy—depends on **not creating such tools**, regardless of current political circumstances or confidence in current influence.

The Commission's task is not to assess who has political influence today, but to examine whether the institutional design provides adequate safeguards across all foreseeable political futures.

Why This Submission Is Relevant to the Commission

The Robodebt Royal Commission (2022–2023, a separate inquiry) examined how Robodebt occurred and recommended reforms intended to prevent repetition.

Schedule 2's timing — enacted approximately 28 months after the Robodebt Royal Commission's report in July 2023 — raises serious questions about whether those lessons have been fully incorporated into subsequent legislative design.

Robodebt Royal Commission Recommendation 17.1 established standards: automated systems should provide clear review pathways, plain-language explanations of decision-making processes, and sufficient transparency to enable independent scrutiny. This submission suggests Schedule 2 may not meet those standards.

The Inspector-General of Intelligence and Security (IGIS) **encouraged the PJCIS Committee to consider what level of public visibility may be desirable** regarding automated arrangements, noting the strong public interest in automated decision-making following the Robodebt Royal Commission and referencing Recommendation 17.1's emphasis on review pathways, explanations, and independent scrutiny
(Inspector-General of Intelligence and Security, Submission No. 145 to the PJCIS Inquiry into the *Combatting Antisemitism, Hate and Extremism Bill 2026*).

This submission respectfully suggests that examining whether similar structural patterns may be emerging in responses to antisemitism may be relevant to this Commission's inquiry into social cohesion.

If early scrutiny does not occur, Schedule 2 may operate for extended periods without fully tested safeguards, potentially allowing harms to accumulate before review — a trajectory observed in both offshore detention frameworks and Robodebt.

This submission respectfully submits that early consideration may help avoid repetition of past harms.

MAPPING TO ROYAL COMMISSION TERMS OF REFERENCE

This section identifies how the issues examined in this submission relate to the Royal Commission's Terms of Reference.

The submission approaches antisemitism and social cohesion primarily through analysis of institutional responses, legislative design, and governance safeguards.

Terms of Reference Area	Where Addressed in Submission	Nature of Contribution
Antisemitism in Australia	Executive Overview; Parts I, III, IV	Examination of antisemitism as a real threat; definitional approaches; moral framing in institutional responses; and safeguards relevant to protecting Jewish communities
Social cohesion	Parts II, V, VI, VIII	Analysis of how institutional design, racial frameworks, historical patterns of democratic breakdown, and institutional risk affect community cohesion and social trust
Responses and remedies	Parts X, XIII, XIV	Alternative automation framework demonstrating lawful approach; legislative and administrative recommendations; detailed technical appendix showing feasible safeguards

Scope of Contribution

This submission complements incident-focused evidence by examining systemic patterns in policy responses and proposing alternative governance frameworks intended to strengthen both protection against antisemitism and broader social cohesion.

Where theoretical or historical analysis is included (for example, the use of Hannah Arendt as an interpretive lens), it is offered solely to assist the Commission in understanding potential long-term institutional implications relevant to its mandate.

PART I: WHAT IS ANTISEMITISM? THE DEFINITIONAL CONTEST

Introduction: The Problem of Contested Definitions

The Antisemitism Bill 2026 purports to combat antisemitism but does not provide a precise statutory definition of the conduct it seeks to regulate. Instead, the legislation references the International Holocaust Remembrance Alliance (IHRA) working definition “for guidance” while establishing significant enforcement mechanisms, including ministerial banning powers and automated intelligence assessment frameworks.

This raises a foundational rule-of-law question: **when the state cannot clearly define prohibited conduct, enforcement risks becoming discretionary rather than predictable.**

The challenge is not merely drafting technique. In contemporary scholarship, policy debates, and public discourse, antisemitism is subject to multiple competing interpretive frameworks. These frameworks diverge sharply on whether criticism of Israel, Zionism, or Palestinian advocacy constitutes antisemitism.

This Part outlines the major definitional frameworks currently in circulation and explains why reliance on contested guidance creates institutional conditions for uncertainty and potential arbitrariness.

The Contemporary Definitional Landscape

At least four prominent frameworks shape current debates about antisemitism. They share consensus regarding classical antisemitism — conspiracy theories about Jews, discrimination against Jews as Jews, or violence targeting Jewish communities — but diverge significantly regarding political speech about Israel and Palestine.

1. IHRA Working Definition (2016)

The IHRA working definition describes antisemitism as:

> “a certain perception of Jews, which may be expressed as hatred toward Jews.”

The definition is accompanied by illustrative examples, several of which concern discourse relating to Israel. These include references to “double standards,” comparisons between Israeli policies and historical regimes, or arguments that the existence of Israel is racist.

The IHRA document states that it is a **non-legally binding working definition** intended for guidance rather than direct application as criminal law.

Key characteristics include:

- Inclusion of examples linking certain criticisms of Israel to antisemitism depending on context;
- Reliance on interpretive judgement about intent and framing;
- Absence of a precise legal threshold for enforcement.

---

2. Jerusalem Declaration on Antisemitism (JDA) (2021)

Developed by scholars of antisemitism and Holocaust studies, the Jerusalem Declaration defines antisemitism as discrimination or hostility toward Jews as Jews.

Unlike IHRA, it explicitly states that the following are not inherently antisemitic:

- advocacy for Palestinian rights;
- criticism of Zionism as a political ideology;
- evidence-based comparisons of Israeli policies to other historical or political systems;
- boycott or sanctions campaigns as forms of political expression.

The JDA therefore draws a sharper distinction between Jewish identity and Israeli state policy.

---

3. Nexus Document (2021)

The Nexus framework provides a legal-analytic approach emphasizing context and intent. It asserts that criticism of Israel or Zionism is not inherently antisemitic unless it relies on antisemitic stereotypes or treats Jews as Jews differently.

Its emphasis lies in distinguishing political speech from prejudice, making it particularly influential in academic and legal policy discussions.

---

4. Project Esther (2024)

A fourth framework — often referenced in political debates — is **Project Esther**, developed within a policy advocacy context. This framework treats antizionism and certain forms of Palestinian advocacy as inherently antisemitic and recommends expanded state responses, including surveillance and institutional restrictions.

Regardless of whether policymakers endorse its conclusions, Project Esther is relevant analytically because it illustrates how the concept of antisemitism may be interpreted expansively in political discourse. References to this framework have appeared in public rhetoric surrounding protests and activism, and its inclusion here assists the Commission in understanding the broader political environment in which accusations of antisemitism arise.

Areas of Consensus and Divergence

Across these frameworks there is broad agreement that the following conduct constitutes antisemitism:

- Holocaust denial;
- violence or threats against Jewish individuals;
- conspiracy theories about Jewish control or collective guilt.

However, the frameworks diverge significantly regarding:

- boycott movements;
- apartheid comparisons;
- antizionism;
- advocacy for Palestinian political rights.

This divergence is not minor. The same statement may be viewed as protected political speech under one framework and antisemitic under another.

The Institutional Challenge: When Definitions Compete

Where a legal regime references “guidance” drawn from contested frameworks, enforcement outcomes may depend on interpretive choices rather than clear legislative standards.

This creates several risks:

1. **Uncertainty for citizens.** Individuals cannot reliably predict whether speech or association may attract sanction.
2. **Delegation of interpretive authority.** Enforcement bodies may rely on external advocacy organisations or prevailing political narratives to assess whether conduct is antisemitic.
3. **Inconsistent application.** Different officials may apply different frameworks depending on context or pressure.

The concern is institutional rather than personal. The issue is not which organisation or viewpoint is correct, but that multiple competing definitions exist simultaneously.

The “Reasonable Member of the Target Group” Standard

Recent anti-hate-speech legislation introduces a test whereby conduct may be assessed by reference to whether a “reasonable member of the target group” would regard it as hateful or antisemitic.

While intended to acknowledge community experience, such a formulation raises rule-of-law questions:

- Who constitutes a “reasonable member”?
- How is reasonableness assessed where communities themselves hold divergent political views?

- Which voices or organisations become de facto interpreters of group perception?

In practice, this structure risks shifting legal judgement toward perception-based assessments that may vary depending on which representatives are treated as authoritative. Where antisemitism itself is subject to competing definitions — IHRA, JDA, Nexus, and frameworks such as Project Esther — reliance on hypothetical community perception may embed one political interpretation without explicit legislative endorsement.

The concern is therefore structural. When criminal or security consequences hinge on subjective group perception rather than clearly defined legal criteria, decision-making authority may effectively migrate away from Parliament and courts toward informal interpretive actors. This creates unpredictability and may chill legitimate political expression.

Evidence of Divergent Interpretations in Australian Public Debate

Public statements by political leaders across multiple jurisdictions demonstrate how differently antisemitism can be interpreted in practice. Statements responding to protests or political advocacy have, at times, treated criticism of Israeli government actions as inherently antisemitic, while other commentators — including Jewish scholars and organisations — reject that equivalence.

This divergence reflects the broader definitional contest described above. It does not imply improper influence by any particular individual or group; rather, it illustrates the absence of a stable consensus capable of supporting criminal or security enforcement.

Implications for the Antisemitism Bill

By referencing IHRA “for guidance” while leaving enforcement mechanisms broad and discretionary, the legislation operates within an environment of definitional ambiguity.

The practical consequences include:

- enforcement outcomes potentially shaped by prevailing political interpretations rather than statutory clarity;
- increased reliance on ministerial judgement;
- a risk that political speech becomes subject to unpredictable regulatory responses.

Where automated intelligence assessments or organisational bans depend on contested definitions, the absence of clear legal standards may undermine confidence in fairness and consistency.

Rule-of-Law Considerations

The principle that criminal laws must provide clear notice of prohibited conduct is well established in liberal democratic systems. When multiple credible frameworks interpret the same speech differently, reliance on non-binding “guidance” risks falling short of that standard.

The issue is not whether antisemitism exists — it clearly does — but whether contested political categories can safely function as triggers for coercive state power.

Conclusion

The definitional contest surrounding antisemitism is not an abstract academic debate. It reflects genuine disagreement among scholars, policymakers, and communities about the boundary between prejudice and political criticism.

By creating enforcement mechanisms that rely on interpretive guidance drawn from contested frameworks — and by incorporating perception-based standards such as the “reasonable member of the target group” test — the legislative framework risks introducing uncertainty into an area requiring the highest degree of clarity.

For rule-of-law purposes, criminal or security measures should rest on clearly defined conduct rather than evolving political interpretations. The next Part considers whether antisemitism, however defined, constitutes a threat requiring state response, and what form that response should take.

PART II

Structural Design, Escalation Risk, and Institutional Safeguards

---

Overview: From Social Harm to Institutional Architecture

Acts of violence directed toward members of the Jewish community have generated understandable fear and urgency. Antisemitism is real and historically consequential. Legislative responses to such harm must be taken seriously.

However, constitutional analysis requires attention not only to the intention of protective measures, but to their structural design. The durability of liberal democratic safeguards depends on institutional architecture that remains stable under political turnover and resistant to over-concentration of discretionary authority.

This Part examines how protective responses interact with institutional design.

I. Violence, Attribution, and Legislative Reaction

A. Complexity of Public Violence

Acts of public violence in Australia over recent decades have arisen from varied and complex causes. Investigations into incidents such as the Port Arthur massacre (1996), the Lindt Café siege in Sydney (2014), the Wieambilla shootings in Queensland (2022), the Bondi Junction stabbings (April 2024), and the Bondi Beach attack (December 2025) demonstrate that high-profile acts of violence in Australia arise from heterogeneous causes, including mixtures of ideological signalling, personal grievance, mental illness, social isolation, or conspiratorial belief systems. The pattern across decades is one of heterogeneity rather than uniform causation.

In several cases, perpetrators were later found to be experiencing significant mental distress, isolation, or delusional belief systems at the time of the offence. In other cases, trials and investigations remain ongoing, and definitive findings regarding motive have not yet been made.

It is important that legislative responses to acts of violence are grounded in established findings rather than provisional attribution. Where events are still subject to criminal process, caution is warranted before drawing broad structural conclusions about causation or social prevalence.

This is not to diminish the seriousness of antisemitism or the fear experienced by Jewish Australians following traumatic events. Antisemitism has deep historical roots and its consequences, when left unchecked, have been catastrophic. The Holocaust stands as the most extreme historical example of state-enabled and industrialised persecution, in which administrative machinery was used to identify, isolate, and eliminate entire populations.

However, the historical lesson of the Holocaust is not only that hatred must be resisted. It is also that the consolidation of unchecked administrative power can become the vehicle through which hatred is operationalised. The machinery of persecution required systems of classification, record-keeping, discretionary enforcement, and institutional insulation from review.

In contemporary Australia, the scale and structure of antisemitic activity differ fundamentally from the conditions that preceded state-sponsored extermination regimes. Australia operates within a constitutional democracy with judicial supervision, parliamentary accountability, and regular political turnover. At the same time, the durability of these safeguards cannot be assumed to be automatic or immune from erosion.

For that reason, any expansion of security powers in response to acts of violence should be evaluated not only by reference to present intentions, but by reference to institutional design. A central structural question arises:

Would the proposed mechanism remain safe if exercised by a future government holding markedly different ideological commitments?

Liberal constitutional traditions are built on precisely this principle. Extraordinary powers should not depend on the presumed virtue of current office-holders, but on structural limits that remain effective under conditions of political change.

Note that Australia is experiencing visible political realignment and debate about constitutional arrangements. Polling for One Nation is surging, disrupting the traditional party balance. Public debate about the royal family may be crystallising into dissatisfaction with the institution of the Crown, and the option of a transition to republic is regularly raised. The stability of our constitutional and political systems cannot be taken for granted into the foreseeable future.

No state can guarantee absolute security. Attempts to eliminate all risk may lead to the accumulation of discretionary authority in ways that reduce transparency, contestability, and judicial supervision. In the long term, this may create vulnerabilities of a different kind — not from private actors, but from institutional overreach.

Protective measures must therefore be assessed against two tests simultaneously:

1. Do they meaningfully address the demonstrated source of harm?
2. Do they preserve the constitutional safeguards that prevent the concentration of coercive administrative power?

Measures that involve systematic classification by ethnicity, "race" or religion, expansion of automated decision-making insulated from review, or the normalisation of exceptional enforcement frameworks require particularly careful scrutiny. Even when introduced for protective purposes, such mechanisms may create institutional architectures capable of future misuse under altered political conditions.

The prevention of future atrocities requires vigilance not only against hatred in society, but also against structural designs that reduce contestability within the state itself. A secure society is not one that eliminates all risk, but one that maintains robust constraints on power even in moments of fear and urgency.

B. The Bondi Incident: Limits of Present Knowledge

The Bondi Beach attack remains subject to criminal proceedings. Full evidentiary findings regarding causation, motivation, and operational dynamics are not yet available in the public domain. The Royal Commission will not possess the complete trial record at this stage.

Permanent structural changes to administrative or security frameworks should therefore be approached with care where the factual basis for systemic conclusions remains incomplete.

C. The Risk of Reactive Institutional Expansion

Moments of trauma can narrow deliberative space. Public pressure for decisive action may accelerate the adoption of exceptional measures. Constitutional safeguards are most vulnerable when urgency compresses scrutiny.

The central question is not whether protection is warranted. It is whether the design of protective measures preserves equal civic status for all Australians and contestability.

II. Protective Measures and the Risk of Structural Escalation

Periods of acute insecurity often generate proposals framed as necessary protective responses. These proposals may arise from genuine community fear and may be motivated by a desire to prevent further harm. However, constitutional analysis requires attention not only to immediate protective intent, but also to the longer-term structural implications of institutional design choices.

Certain categories of protective measure warrant particular caution.

First, proposals involving identity-based spatial arrangements — such as exclusive security perimeters, restricted-access zones, or the formal separation of communities for protective purposes — risk normalising differential treatment along ethnic, "racial" or religious lines. Even when introduced voluntarily or under benevolent conditions, such arrangements may embed classification practices into administrative systems. Over time, distinctions framed as protective can alter expectations about equal civic status and access to public space.

Modern liberal democracies have historically resisted formalised spatial differentiation by identity precisely because systems of classification combined with territorial demarcation can become structurally vulnerable under conditions of political change. The relevant question is not whether such measures are presently well-intentioned, but whether they preserve the principle that all citizens participate in a common civic order governed by equality before the law.

Second, proposals to authorise or normalise armed non-state security groups — even when coordinated with police — raise complex constitutional considerations. A central feature of constitutional democracy is the state’s monopoly on the legitimate use of force, exercised through accountable public institutions. The emergence of parallel communal enforcement structures may

unintentionally fragment that principle. Even where initially conceived as defensive, such arrangements can contribute to the perception that security is organised along communal rather than civic lines.

It also introduces the possibility of reciprocal arming dynamics. Where one community is visibly authorised to maintain armed protective capacity, other groups may perceive a need to adopt similar measures for their own protection, whether formally sanctioned or otherwise. Even isolated misunderstandings or low-probability incidents in such an environment may escalate more rapidly once multiple armed actors are present. In more extreme circumstances, such dynamics can contribute to cycles of reactive enforcement and counter-enforcement that are difficult to contain. Historical experience demonstrates that fragmented coercive authority between rival armed groups can lead to rapid and unintended escalation in public settings. Liberal democratic systems therefore centralise the legitimate use of force within accountable public institutions in order to avoid such risks.

This structural concern is not directed at any specific community. It reflects a broader constitutional principle: protective measures should reinforce equal civic status and centralised accountability, not embed identity-based differentiation into enforcement architecture.

Historical experience demonstrates that systems of classification, separation, and delegated coercive authority, once institutionalised, can outlast the conditions that produced them. Mechanisms introduced in response to immediate fear may, over time, alter the relationship between communities and the state in ways that are difficult to reverse.

Accordingly, proposed protective responses should be assessed against three institutional tests:

1. Do they preserve equal civic status across communities?
2. Do they maintain the state's neutral and accountable role as guarantor of security?
3. Do they avoid embedding identity-based classification into spatial or enforcement structures?

Security measures that satisfy these criteria are more likely to strengthen both safety and social cohesion. Measures that entrench separation or parallel enforcement capacity, even if introduced with protective intent, risk producing longer-term structural fragility.

III. Administrative Classification and the "Right to Have Rights"

Hannah Arendt's analysis of statelessness and administrative governance in *The Origins of Totalitarianism* examined how legal vulnerability can arise not only from overt violence, but from the gradual reorganisation of political belonging. Central to her account was what she termed the "right to have rights" — the condition of being recognised within a political community capable of enforcing legal claims.

Arendt argued that the most consequential shifts in political protection often occur at the level of administrative design. When individuals move from juridical status into bureaucratic management

— particularly through classification mechanisms that determine belonging, exclusion, or risk — their capacity to contest state action may be reduced. Her concern was not solely with the existence of prejudice, but with the institutional conditions that allow administrative discretion to operate with limited review.

Her later work, including *Eichmann in Jerusalem*, was controversial and remains debated. Some readers objected to aspects of her interpretation of institutional behaviour during the Holocaust. That debate is important and ongoing. Nonetheless, her structural insight — that administrative systems can reshape the conditions of political protection — remains influential in constitutional and political theory.

The relevance of this analysis in the present context lies not in direct historical comparison, but in institutional design. Where security frameworks incorporate identity-based classification, automated assessment, or reduced contestability, careful scrutiny is warranted. The constitutional question is whether such mechanisms preserve the equal civic status of individuals and maintain effective avenues for review.

The lesson drawn here is limited and structural: protective intentions do not eliminate the need for institutional safeguards. Systems that categorise, manage, or assess individuals must remain embedded within a framework that preserves contestability, equal standing, and judicial supervision. The prevention of future injustice depends not only on resisting hatred, but on maintaining the conditions under which all persons retain the practical “right to have rights.”

IV. Automation, Contestability, and Administrative Insulation

Automated assessment frameworks and ministerially expandable instruments introduce additional structural considerations.

Where decision-making processes are automated, notification reduced, or review narrowed, the practical ability of individuals to challenge state action may be diminished. Administrative efficiency must be balanced against contestability.

The constitutional concern is not automation per se. It is insulation from review. Systems that reduce transparency, centralise discretion, or narrow judicial supervision alter the equilibrium between security and accountability.

The structural test remains consistent: does the mechanism preserve equal standing and effective avenues of review?

V. Definitions, Moral Language, and Institutional Design

Definitions such as the IHRA working definition of antisemitism play an important interpretive role. Moral language reflects the historical gravity of antisemitism and the catastrophic consequences of neglect.

However, when moral narratives inform legislative architecture, care must be taken to ensure that urgency does not justify the expansion of exceptional administrative mechanisms without adequate structural limits.

Section 75(v) of the Constitution reflects a foundational safeguard: executive action remains subject to judicial supervision. This principle reinforces the idea that extraordinary powers should not depend upon the virtue of present office-holders, but upon structural constraints that endure under political change.

VI. The Security–Liberty Trade-Off

This submission recognises that governments face genuine pressure to respond decisively to threats against vulnerable communities. Protective measures are often adopted in good faith and with the intention of preventing harm. The structural concern raised here is not directed at those intentions, but at the long-term institutional consequences that can arise when emergency responses become embedded as permanent legal frameworks.

No state can guarantee absolute security. Attempts to eliminate all risk may lead to the accumulation of discretionary authority in ways that reduce transparency and contestability.

The long-term preservation of social cohesion depends not only on preventing harm, but on maintaining institutional limits that prevent concentration of coercive power.

Durable systems protect both safety and liberty by preserving review, equal status, and accountability.

VII. Political Volatility and Institutional Durability

Over recent decades, Australian public debate on national security and social cohesion appears to have moved toward a narrower range of policy positions. Measures that might previously have attracted significant ideological contestation increasingly receive bipartisan support. While this convergence may reflect shared assessments of risk following major incidents of violence, it may also reduce the adversarial scrutiny that traditionally functions as a safeguard within liberal democratic systems.

Periods of heightened public anxiety often generate expectations that governments will demonstrate visible and immediate action. Responses such as increased police presence, expanded surveillance powers, or the consideration of additional security structures can therefore acquire strong symbolic importance. These measures may reassure communities in the short term; however, they also raise structural questions about proportionality, oversight, and long-term institutional design.

Public discussion following the Bondi attack illustrates this dynamic. Proposals involving enhanced security capabilities, expanded police equipment, or new forms of community-based protection demonstrate the powerful social demand for safety after traumatic events. Yet liberal constitutional traditions emphasise that responses to crisis must remain consistent with enduring principles of

accountability, contestability, and restraint. The objective is not to deny the reality of threats, but to ensure that institutional responses do not unintentionally create frameworks that could later be repurposed in ways that undermine social cohesion.

This submission therefore approaches recent legislative developments within a broader institutional context rather than a partisan one. The concern is not the intentions of any particular political actor, but the possibility that a sustained climate of security consensus may reduce the diversity of perspectives ordinarily present in parliamentary debate. When scrutiny becomes limited, complex measures — particularly those involving automated decision-making, secrecy provisions, or executive discretion — may be introduced without full appreciation of their structural implications.

A liberal democratic system recognises an inherent tension between security and freedom. Absolute security cannot be guaranteed; attempts to eliminate all risk may require expanding state power in ways that introduce new vulnerabilities of their own. The strength of democratic institutions lies not in their ability to promise complete safety, but in their capacity to maintain open contestation and independent oversight even during periods of fear or uncertainty. Understanding this tension provides essential context for evaluating Schedule 2 changes to the ASIO Act in the Combatting Antisemitism, Hate and Extremism Bill 2026 and similar frameworks.

Part II Conclusion

Protective responses to antisemitism must be grounded in equal civic status, judicial supervision, and institutional contestability.

The central structural question is constant:

Would this mechanism remain safe if exercised by a future government holding markedly different commitments?

Institutional safeguards must remain effective under conditions of constitutional change, political realignment, or executive turnover. Mechanisms that depend upon the stability of present arrangements, or the assumed moderation of future office-holders, are inherently fragile.

The preservation of the “right to have rights” requires that security measures reinforce — rather than weaken — the constitutional conditions under which all persons remain equal before the law.

Part III: Antisemitism Is Real

Introduction: Acknowledging the Reality of Harm

Before examining the structure and potential consequences of the Antisemitism Bill, this submission begins from a clear premise:

Antisemitism is real, harmful, and deserving of serious response.

Jewish Australians have reported increased fear for their safety.
Synagogues and community events have adopted enhanced security measures.
Incidents involving antisemitic language, vandalism, and threats have been documented across Australia.

These concerns are not abstract. They reflect lived experiences within a community that carries a long historical memory of persecution.

Nothing in this submission seeks to minimise those experiences.
The question addressed here is not whether antisemitism requires action — it is **what form that action should take**, and whether the proposed legislative framework is likely to strengthen or weaken long-term social cohesion and safety.

Contemporary Context

Far-Right Extremism

Across democratic societies, analysts have observed a resurgence of far-right movements characterised by:

- ethnonationalist narratives,
- conspiracy ideologies,
- hostility toward minorities,
- and the strategic use of online networks for mobilisation.

In Australia, several groups and informal networks associated with neo-Nazi or white supremacist ideology have drawn attention from law-enforcement agencies and researchers.
These movements have historically incorporated antisemitic themes alongside anti-Muslim and anti-immigrant rhetoric.

The presence of such movements forms an important part of the contemporary context in which Jewish Australians assess their security.

Recorded Incidents

Public reporting indicates a rise in antisemitic incidents including:

- graffiti or vandalism targeting Jewish institutions,
- harassment directed at individuals perceived to be Jewish,
- online abuse,
- and threats requiring police attention.

Each incident contributes to an atmosphere of vulnerability, particularly when combined with global events that heighten communal tension.

Recognising this reality is essential for any constructive policy discussion.

Public Demonstrations Since October 2023

Since the escalation of violence in Israel and Gaza in October 2023, Australia has seen sustained public demonstrations reflecting strong and often opposing views.

Two broad patterns have emerged.

Large-Scale Gaza Solidarity Demonstrations

Across multiple cities, large gatherings have been organised calling for ceasefires, humanitarian relief, or broader political change.
Many of these demonstrations have involved families, students, and community organisations, and have been conducted with permits and police coordination.

While controversial slogans or isolated incidents have attracted media attention, the majority of events appear to have occurred without large-scale violence.

Counter-Demonstrations and Far-Right Presence

Smaller counter-protests have also occurred, including participation by individuals or groups associated with far-right ideology.
In some cases, confrontational dynamics have arisen when opposing groups gather in close proximity.

From a policy perspective, distinguishing between peaceful political advocacy and organised extremist activity is essential.
Failure to maintain this distinction risks conflating democratic protest with genuine security threats.

Instrumentalisation and Narrative Competition

Research examining contemporary far-right strategy suggests that some movements employ pro-Israel or pro-Jewish rhetoric to frame themselves as defenders of Western civilisation, even where underlying antisemitic beliefs remain present.

Analyses such as *The New Philosemitism* (Institute for Strategic Dialogue, 2020) describe how expressions of support for Jewish communities may function strategically — helping far-right actors present themselves as mainstream while advancing anti-Muslim agendas.

This dynamic creates a complex environment in which:

- genuine concern for Jewish safety,
- geopolitical conflict,
- and domestic political debate

can become intertwined in ways that obscure the primary sources of antisemitic risk.

The implication for policymakers is not that support for Jewish communities is insincere, but that legislative responses should be carefully designed to avoid unintentionally amplifying polarisation or misdirecting enforcement priorities.

The Risk of Misdirected Policy Responses

Where public debate frames antisemitism primarily through the lens of geopolitical protest, there is a possibility that legislative measures may focus on highly visible political speech rather than on the organised networks historically associated with antisemitic ideology.

This submission does not assert predetermined enforcement outcomes.
Instead, it raises a structural question for consideration:

Will new powers primarily address violent extremism, or will they be applied in contexts of contested political expression?

Experience from other policy domains suggests that laws introduced during periods of heightened fear can expand beyond their original intent.
Ensuring proportionality and clarity at the design stage is therefore critical.

The Importance of Distinguishing Jews from State Policy

Many Jewish Australians hold diverse views regarding Israeli government actions and the broader Israel-Palestine conflict.
A legislative framework that appears to conflate Jewish identity with support for any particular state policy risks unintentionally reinforcing stereotypes that treat Jews as politically monolithic.

Protecting Jewish Australians from discrimination should therefore be framed in a way that recognises internal diversity within Jewish communities and avoids attributing collective responsibility for international events.

Why Careful Framing Matters

If legislation is perceived as privileging certain groups or suppressing legitimate political debate, it may produce unintended social consequences:

- resentment directed toward minority communities,
- erosion of trust in institutions,
- or the strengthening of narratives used by extremist actors.

Conversely, a framework grounded in universal human rights — protecting individuals from violence and discrimination regardless of identity — may reduce perceptions of hierarchy while still addressing genuine threats.

Conclusion: Affirming Reality While Questioning Method

Antisemitism is a serious and growing concern.
Acknowledging that reality is the starting point for any meaningful policy discussion.

However, effective protection requires careful alignment between **the nature of the threat** and **the tools used to respond to it**.
Legislation that inadvertently conflates political dissent with extremism, or that embeds rigid identity categories into law, may undermine the very social cohesion it seeks to defend.

The analysis that follows therefore proceeds from a position of respect for Jewish safety concerns while examining whether the proposed legislative architecture achieves its stated purpose in a proportionate and sustainable way.

The reality of antisemitism therefore cannot be in doubt.
The question becomes how institutions interpret threat, and how language shapes the scope of policy responses.
The next Part examines the role of moral framing in institutional behaviour.

PART IV: LANGUAGE, MORAL FRAMING, AND INSTITUTIONAL POWER

Introduction: Why Language Matters in Moments of Fear

Having established the reality of antisemitism, this Part explores how institutional language influences the perceived scale of threat and the expansion of administrative power.

This submission has documented how institutional design, definitional ambiguity, and automated decision-making can expand administrative power.
This Part examines a related but distinct issue: **how moral language — particularly the framing of threats as "evil" — can influence institutional behaviour during periods of perceived crisis.**

The response to the Bondi Beach attack demonstrates how rapidly public discourse can shift toward existential framing.
When governments and institutions adopt moralised language, the effect is not merely rhetorical. It can shape how risks are assessed, how powers are justified, and how restraint is perceived.

This analysis does **not** question the seriousness of antisemitism or the legitimacy of efforts to combat discrimination and violence.
Rather, it considers whether particular forms of language — especially when embedded within policy frameworks — unintentionally alter the balance between democratic safeguards and executive urgency.

The IHRA Preamble and the Installation of a Moral Frame

The International Holocaust Remembrance Alliance (IHRA) Working Definition of Antisemitism opens with a preamble referencing the Stockholm Declaration's call to **"fight those evils."**

This submission respectfully observes that the placement of the word *evil* before the definition proper establishes a moral lens through which subsequent interpretation may occur.

Before any conduct is described, before any examples are listed, the framework situates antisemitism within a category of existential moral threat.

Language of this kind can be powerful and emotionally resonant.
However, when incorporated into governance frameworks or institutional policies, it may have structural consequences.

What Moral Framing Does Inside Institutions

When institutions with coercive authority adopt language associated with existential struggle, several subtle shifts may occur.

1. From Conduct to Identity

Rule-of-law systems ordinarily evaluate **actions**:

- unlawful conduct
- discriminatory behaviour
- threatening speech
- criminal activity.

Moral framing can shift attention toward **identity or perceived alignment**, encouraging a view that certain ideologies or associations represent inherent danger rather than specific unlawful acts.

The result is not necessarily intentional bias, but a change in how risk is conceptualised.

---

2. The Elastic Quality of Moral Categories

Once a framework is framed around combating “evil,” its application can expand beyond the original context.

Across modern policy debates, language has often moved from descriptive to existential:

“stop the boats,” “dole bludgers,” “terrorists,” “extremists.”

Placed together in a single line, these phrases sound abrupt — even jarring — compared to the careful legal language used throughout this submission.
That contrast is intentional. It illustrates how quickly discourse can shift from analytical description to moral shorthand.

Historically, such shifts have often preceded expansions in administrative authority.
When institutions perceive an existential threat, measures that would normally be considered exceptional may begin to appear proportionate.

---

3. Perceived Existential Threat and Administrative Expansion

Across migration policy, welfare administration, and national-security discourse, a recurring pattern emerges:

- moral framing elevates perceived urgency;
- urgency reduces tolerance for procedural delay;
- reduced tolerance for delay encourages reliance on automated or discretionary mechanisms.

In these conditions, institutions may begin to treat administrative expansion as proportionate to an existential threat rather than as an exceptional measure requiring careful restraint.

This dynamic does not depend on ideology or intent.
It is a structural response to perceived crisis.

A Thought Experiment: Substituting Moral Language with Legal Language

The Parliament has recently enacted legislation framed in legal rather than moral terms — prohibiting unlawful hate speech and setting thresholds based on harm and reasonableness.

Against that background, it may be useful to consider a rhetorical exercise.

If the IHRA preamble's call to "fight those evils" were rewritten using legal language — for example, "combat criminal or unlawful acts of antisemitism" — the emotional tone would shift markedly.

The purpose of raising this question is not to rewrite IHRA or challenge its historical significance. Rather, it illustrates how language influences institutional perception:

- **"Evil"** invites moral urgency and identity-based framing.
- **"Criminal" or "unlawful"** invites evidentiary standards, proportionality, and judicial oversight.

The substitution demonstrates how framing can alter the balance between moral condemnation and legal analysis.

Language and Institutional Behaviour: Historical Patterns

Across the broader submission, similar dynamics appear in multiple policy contexts.

Migration debates framed around existential threat narratives.
Welfare debates framed through moralised depictions of undeserving recipients.
National-security debates framed around extremism.

When viewed together, these episodes suggest a recurring pattern:

moral framing → perceived existential risk → expansion of administrative authority.

This observation does not imply that earlier policies were identical or motivated by the same considerations.
It simply highlights how language can shape institutional responses long before legal thresholds are examined.

Why Universities and Public Institutions Become Particularly Vulnerable

Universities, public agencies, and cultural institutions often adopt external frameworks without the interpretive tools necessary to distinguish moral rhetoric from legal obligation.

When moral framing becomes embedded in policy guidance:

- administrators may act conservatively to avoid reputational risk;
- contested political speech may be treated as institutional liability;
- identity-based claims may carry greater weight than evidence-based analysis.

The resulting chilling effect can occur even without formal enforcement.

The Liberal-Democratic Alternative: Legal Language and Institutional Restraint

Liberal democracies traditionally respond to harm using legal categories:

- criminal
- unlawful
- discriminatory
- threatening
- harmful.

These terms require evidence, enable review, and preserve proportionality.

Moral language, by contrast, can reduce the perceived legitimacy of procedural safeguards. When institutions frame their role as “fighting evil,” restraint may appear morally suspect rather than constitutionally necessary.

The lesson drawn here is not that moral language should disappear from public life. Rather, that governance frameworks function best when grounded in legal standards rather than existential framing.

Narrative Framing, Motive Attribution, and Institutional Acceleration

Modern democratic politics relies heavily on message testing, focus groups, and polling to refine public communication. Campaign narratives are often shaped not only by policy substance, but by how issues are framed in emotionally resonant terms. This is a normal feature of competitive electoral systems.

Political communication frequently adopts simplified narrative structures that distinguish between protectors and perceived sources of harm. Such framing can reduce complex institutional questions to moral binaries.

Research in political communication suggests that messages become more persuasive when they:

1. Define a moral in-group (“the good community”),
2. Identify a perceived source of harm or threat,
3. Attribute a credible motive to that source, and

4. Position proposed measures as protective of the in-group.

The attribution of motive plays a particularly stabilising role. Messages framed around ignorance, discrimination, greed, extremism, or indifference tend to be more believable to audiences than narratives that describe harm without explanation. Once motive is attached, threat perception often becomes more durable.

Recent Australian public debates have illustrated how moral framing can harden divisions. During the Voice referendum campaign, for example, public discourse at times moved beyond institutional design questions to arguments framed in moral terms — including claims that opposition reflected racism, and counter-claims that such accusations themselves were divisive. The resulting dynamic narrowed deliberative space and intensified identity-based alignment.

The purpose of this observation is not to assess the merits of that referendum. It is to note a structural pattern: when public debate becomes moralised, disagreement can be interpreted as moral failure rather than institutional caution. Once that occurs, proposals framed as protective responses may acquire heightened symbolic urgency.

This dynamic is not unique to any single issue area. In past Australian debates concerning offshore detention, welfare compliance, counter-terrorism, and other high-salience matters, vivid imagery and motive attribution have contributed to shaping public perception of risk. In each case, political incentives have rewarded visible demonstrations of protection.

Such dynamics can generate institutional acceleration. When symbolic protection becomes politically salient, there may be pressure to:

- Expand enforcement capacity,
- Reduce procedural friction,
- Increase executive discretion, or
- Signal resolve through structural change.

The risk is not that protection is sought. The risk is that institutional design becomes driven by narrative urgency rather than long-term structural evaluation.

In moments of heightened moral framing, scrutiny may narrow. Mechanisms that would otherwise receive extended constitutional debate — particularly those involving classification, automation, or expanded discretionary authority — may be assessed primarily through the lens of immediate reassurance rather than durable institutional stability.

The relevance to the present submission is structural rather than partisan. Legislative frameworks adopted during periods of narrative intensity must remain safe under conditions of political change, ideological turnover, and shifting public mood. Institutional design should not depend upon the continuation of present narratives or assumptions about future office-holders.

Understanding the interaction between narrative framing and institutional acceleration is therefore central to evaluating Schedule 2 and related reforms. Durable democratic systems are those that maintain contestability and equal civic status even when political messaging compresses deliberative space.

Connection to the Broader Submission

This analysis complements earlier Parts examining structural risks:

- Part I demonstrated how contested definitions create uncertainty about enforcement boundaries.
- Part II examined how structural design features create escalation risks and require institutional safeguards.
- Part III established that antisemitism is real and requires serious response.

Part IV adds an additional layer: the role of language in shaping how institutional structures are interpreted and applied, and how moral framing can amplify urgency while reducing procedural restraint.

Transitional Bridge: From Language to Institutional Design

The analysis in this Part is not an abstract discussion of rhetoric.
Across this submission, a consistent pattern has emerged: **language shapes institutional behaviour, and institutional behaviour shapes the exercise of state power**.

Earlier Parts established that antisemitism is real and requires serious response (Part III), while showing how contested definitions create enforcement uncertainty (Part I) and how racial categorisation frameworks raise questions about essentialism and equality (Part V).

Part IV adds a further dimension — the way moral framing, particularly language invoking existential threat or "evil," can amplify institutional urgency and reduce tolerance for procedural restraint.

Seen together, these elements form a single analytical chain:

- **Moral framing** elevates perceived threat;
- **Definitional ambiguity** expands interpretive discretion;
- **Administrative mechanisms** operationalise that discretion at scale.

The concern raised by this submission is therefore structural rather than partisan.
When these elements converge, institutions may expand coercive powers not through explicit intention, but through the cumulative effects of urgency, uncertainty, and technological capacity.

The following sections examine how these dynamics have manifested in Australia's policy trajectory—from offshore detention through Robodebt to Schedule 2—and propose safeguards ensuring that responses to antisemitism remain effective while preserving the rule-of-law principles that protect all communities.

Part V: The Racial Framework Problem

Introduction: When Protection Relies on Fixed Identity

Part I demonstrated that contested definitions of antisemitism may create risks of inconsistent or arbitrary enforcement.

Building on the discussion of institutional framing in Part IV, this Part turns to a related structural question: how the Bill's reliance on racial categorisation shapes the way legal protection is constructed and applied.

The analysis does not question the need to protect Jewish Australians from genuine threats. Rather, it examines whether anchoring protection to fixed racial or ethnocultural categories introduces unintended tensions within a legal system committed to equality before the law.

Evidence presented to the Parliamentary Joint Committee on Intelligence and Security by the Executive Council of Australian Jewry in a public hearing on the 14th of January indicated that Jewish identity was discussed primarily in ethnocultural or racial terms for the purposes of protection under the legislation. In contrast, other communities — including Muslims, Christians, Buddhists and Hindus — were characterised predominantly as faith communities rather than racial groups.

This distinction raises broader questions about how identity is defined in law, and whether frameworks grounded in fixed categories adequately reflect the fluid, plural and evolving nature of human identity.

Can racism be addressed effectively by re-asserting racial categories themselves?

The analysis that follows does not question the sincerity of those advocating stronger protections for Jewish Australians. Antisemitism is real and harmful, and safeguarding Jewish communities is a legitimate objective of public policy. The issue examined here is structural rather than personal: whether grounding legal protection in essentialised identity categories may risk reproducing the same classificatory logic that historically enabled discrimination, even when deployed with protective intent.

Identity as Fluid Rather Than Fixed

Throughout the Committee evidence, it was also acknowledged that Jewish identity is diverse — religiously, culturally, politically, and geographically. This diversity reflects a broader reality of human identity.

Human identity is not static.
Each person is born of two parents, themselves born of two parents, in a recursive pattern extending across centuries of migration, conversion, language change, and cultural exchange. Jewish communities, like Muslim, Christian, Hindu, or secular communities, encompass religious

believers, cultural participants, converts, intermarried families, and individuals whose relationship to tradition evolves across a lifetime.

From this perspective, identity is:

- plural rather than singular,
- relational rather than essential,
- historically shaped rather than biologically fixed.

Legal frameworks that attempt to stabilise identity into immutable categories risk oversimplifying this complexity. When protection depends on proving that a group constitutes a “race,” the law may unintentionally privilege rigidity over lived human diversity.

Racism as a Colonial Concept

Modern racial classification did not emerge naturally from human difference.
It developed historically alongside colonial governance and nineteenth-century racial science, which sought to categorise populations for purposes of administration, hierarchy, and control. Colonial systems transformed fluid cultural identities into fixed legal labels — defining who could travel, own land, vote, or belong.

Contemporary scholarship in anthropology, genetics, and international human rights law broadly recognises race as a social construction rather than a biological reality. The United Nations and numerous scientific bodies have emphasised that national, religious, and cultural communities do not correspond to distinct biological races.

This historical context creates a paradox.
If racism originates in rigid classification, relying on those same classifications as the basis of legal protection may unintentionally perpetuate the logic of colonial hierarchy — even when deployed for benevolent purposes.

The Race–Religion Distinction in Evidence to the Committee

Evidence before the Committee suggested that:

- Jews may be treated as an ethnocultural or racial group;
- Muslims and other religious communities may be treated primarily as faith communities.

The distinction appears to rely on the idea that Jewish identity is inherent or immutable, whereas other religious identities are chosen beliefs. Yet the lived reality of communities challenges this binary.

Children raised in Muslim, Christian, Jewish, or other traditions typically inherit language, culture, and community from birth. Converts enter each tradition through recognised processes. Individuals may move toward or away from religious practice throughout life. In this sense, no major religious community fits neatly into a purely voluntary or purely immutable category.

If Jewish identity is treated as racial because it includes ancestry, then similar reasoning could apply to many ethnoreligious communities worldwide. The distinction therefore risks appearing inconsistent or strategically applied rather than conceptually stable.

Descent-Based Identity Frameworks Across Communities

The Ethnocultural Categorisation in Law

In testimony before the Parliamentary Joint Committee on Intelligence and Security on 14 January 2026, Peter Wertheim of the Executive Council of Australian Jewry explained the legal basis for treating Jewish identity as "ethnocultural" rather than purely religious:

> "The word 'race' in the context of antidiscrimination legislation has a very different meaning. It means a group that has a shared past, a shared history and shared culture, language and other ethnocultural identity that is accepted, both within the group and more generally by other people in society. In that sense, it has long been established in the jurisprudence of Australia, New Zealand, Canada, the United Kingdom and the United States that Jews are an ethnocultural and ethnoreligious community in that broader sense and are therefore protected under the rubric of the word 'race' in antidiscrimination legislation."

When asked about converts to Judaism, Wertheim confirmed they receive the same protection: "If you convert to Judaism through a religious process, then you become part of the community in its totality and in every sense. It would encompass the ethnocultural, the linguistic, the Hebrew language, the Hebrew culture—everything that goes with it would be part of that."

This ethnocultural categorisation may derive in part from traditional Jewish law (halakha), which defines Jewish identity matrilineally—a person is considered Jewish if their mother is Jewish. This descent-based framework creates a form of inherited identity that supports categorisation as an ethnic group rather than purely a religious community.

However, similar descent-based frameworks exist in other religious traditions. Islamic law traditionally considers children of a Muslim father to be Muslim through patrilineal descent. Children raised in Muslim households typically inherit language, cultural practices, and community membership from birth, much as children raised in Jewish households do.

Both frameworks include provisions for conversion—Jewish conversion processes exist, as do Islamic conversion practices (shahada). Both communities include individuals who move toward or away from religious observance across their lifetimes. Both encompass profound internal diversity—ethnic, linguistic, theological, and political.

If descent-based inheritance of identity justifies treating one community as ethnocultural/racial, the logical consistency of that principle requires similar treatment of other communities with comparable descent frameworks. The differential categorisation—Jews as racial, Muslims as religious—appears grounded more in legal strategy than in coherent conceptual distinctions.

The Practical Legal Consequence: Differential Protection

This categorisation has practical legal consequences within Australian law. Under existing Commonwealth legislation, the Racial Discrimination Act 1975 (particularly section 18C) prohibits conduct that offends, insults, humiliates, or intimidates based on race, colour, or national or ethnic origin. No equivalent Commonwealth legislation provides comparable protection based on religion.

The exposure draft of the Combatting Antisemitism, Hate and Extremism Bill 2026 proposed a criminal racial vilification offence covering "race, colour or national or ethnic origin" but not religion. While this specific provision was removed from the final legislation passed in January 2026, the broader framework within Australian law treats racial and religious vilification differently.

Multiple witnesses before the Parliamentary Joint Committee on Intelligence and Security highlighted this differential treatment:

Equality Australia (Heather Corkhill, Legal Director): "The new offence of 'promoting or inciting hatred' and the aggravated sentencing provisions apply only to race, leaving our community without equivalent protections from hate and violence... The very criminal laws being amended by this bill already recognise LGBTIQ+ people as at-risk communities alongside women, people of faith and people with a disability. We fought for this inclusion federally just a year ago... To now split and fracture these protections only one year after recognising our communities at a federal level is a position that Jewish leaders have rejected."

Australian National Imams Council (Imam Shadi Alsuleiman, President): "Islamophobia is not recognised as a protected category despite well-documented increases in Islamophobia and anti-Muslim harm and hate in Australia... We understand that race is protected but attributes to faith are not, and Islamophobia's been on the rise."

Australian Human Rights Commission (Hugh de Kretser, President): "We should protect all communities that face hate crimes. This means ensuring protections cover not only race related hate crimes but also hate crimes against people on the grounds of religion, sex, sexual orientation, gender identity and disability."

The practical effect of categorising Jewish identity as racial and Muslim identity as religious is that:

- Conduct targeting Jews may attract protection under racial vilification frameworks (Commonwealth and state-based);
- Comparable conduct targeting Muslims may not attract equivalent protection if framed as religious rather than racial vilification;
- This creates a hierarchy of legal protection based on how communities are categorised rather than on the nature or severity of the vilifying conduct.

This submission does not argue that Jewish Australians do not deserve protection. Rather, it suggests that constructing protection through racial categorisation—particularly where similar descent-based frameworks exist in other communities but are not afforded equivalent categorisation—may create inconsistencies that undermine both effectiveness and perceived fairness.

A legal framework grounded in protecting **all individuals from discrimination, hatred, or violence** motivated by identity (whether racial, religious, ethnic, or national) may provide more consistent and durable protection than one that depends on proving which communities constitute "races."

The Problem of Legal Essentialism

Treating identity as immutable can produce unintended consequences.
Descent-based frameworks — whether religious, legal, or cultural — may appear to reinforce arguments that identity is fixed by lineage. Yet historically, descent-based classification has also been used to exclude, marginalise, or persecute.

Legal history demonstrates that once identity categories are embedded in law, they tend to acquire administrative momentum. Courts cite precedent, bureaucracies apply classifications, and political debates become structured around group hierarchy. Even when introduced to protect vulnerable communities, such frameworks can be reinterpreted by future governments in ways never originally intended.

The concern raised here is therefore structural:
When law defines people through rigid categories, it risks freezing dynamic human identities into permanent administrative labels.

Identity Diversity Across Communities

Jewish communities themselves demonstrate profound diversity:

- religiously (Orthodox, Reform, secular, cultural),
- politically (Zionist, non-Zionist, anti-Zionist),
- ethnically (Ashkenazi, Sephardi, Mizrahi, Ethiopian, and others),
- geographically and linguistically.

Muslim and Christian communities display equally complex diversity across ethnicity, theology, language, and political outlook.

Recognising this shared complexity helps move the conversation away from hierarchy toward universality.
Rather than asking which communities constitute “races,” a liberal democratic framework may instead focus on protecting individuals from discrimination motivated by race, ethnicity, or religion — without requiring the state to declare any group inherently racial.

Human Rights Without Racial Hierarchy

International human rights law offers an alternative model.
Treaties such as the International Convention on the Elimination of All Forms of Racial Discrimination prohibit conduct motivated by race, colour, descent, or ethnic origin without requiring the creation of fixed racial groups within domestic law.

This approach shifts the focus from identity classification to behaviour:

- Violence, harassment, or discrimination motivated by hatred is prohibited regardless of who the victim is.
- Protection attaches to the individual rather than to a legally privileged category.
- Equality before the law is preserved while still recognising the reality of antisemitism, Islamophobia, and other forms of prejudice.

Such a framework aligns with the broader principle that liberal democracies safeguard diversity by resisting essentialism.

The Risk of Re-creating Colonial Logic

The historical record shows that racial categories introduced for protective or administrative purposes often expand beyond their original intent. Once embedded, they may be reinterpreted by different political actors or applied to new groups.

This is not an argument that present policymakers intend harm.
It is an observation about institutional dynamics: systems built on fixed identity categories can evolve in ways that reinforce hierarchy rather than equality.

If identity is inherently fluid — shaped by migration, intermarriage, belief, and lived experience — then legal frameworks premised on immutability risk narrowing human reality to fit administrative convenience.

Conclusion: Protection Without Essentialism

The central insight of this Part is not that Jewish identity lacks depth or historical continuity.
Rather, it is that **all human identities are complex, layered, and evolving**.

Attempting to secure protection by framing any community as racially immutable may inadvertently reproduce the classificatory logic from which modern racism itself emerged. Liberal democratic law can respond more consistently by protecting individuals from discriminatory conduct without requiring the state to define which communities constitute races.

Such an approach:

- affirms the reality of antisemitism and the need for protection,
- recognises similar vulnerabilities faced by other communities,

- and preserves the principle that identity is plural, dynamic, and human rather than fixed by legal definition.

PART VI: INSTITUTIONAL DESIGN AND HISTORICAL LESSONS

Introduction: Why Machinery Matters More Than Motive

Parts I through V established several foundations:

- Definitions of antisemitism remain contested across different frameworks
- Antisemitism presents genuine threats requiring policy responses
- Moral framing influences how institutions perceive and respond to threats
- Identity-based legal frameworks raise questions about essentialism and equality

This Part examines a related question: **What role does institutional architecture play in determining whether protective responses strengthen or weaken democratic safeguards?**

The analysis draws on political theory, particularly Hannah Arendt's scholarship on totalitarianism, to consider how administrative systems—regardless of their stated protective purposes—can develop characteristics that enable systematic harm when certain design features converge.

This inquiry is not abstract. Recent Australian experience, including the Robodebt scheme (2016–2019), demonstrates how automated decision-making systems operating with limited oversight can produce widespread harm even when designed by officials acting within what they understood to be their lawful authority.

The relevance to this Commission lies in understanding whether institutional safeguards proposed for responding to antisemitism incorporate lessons from documented administrative failures, or whether certain structural patterns may recur.

Hannah Arendt: Historical Context and Analytical Framework

Background and Scholarly Contribution

Hannah Arendt (1906–1975) was a German-Jewish political philosopher whose work emerged from direct experience of Nazi persecution and statelessness.

Biographical context:

- Fled Nazi Germany in 1933
- Stateless refugee for 18 years
- Interned in France before escaping to the United States in 1941
- Observed the collapse of democratic institutions and the rise of totalitarian systems

Major works relevant to this analysis:

- *The Origins of Totalitarianism* (1951)

- *Eichmann in Jerusalem: A Report on the Banality of Evil* (1963)

Arendt's scholarship is widely recognised in political science, philosophy, and legal theory. Her analysis of how ordinary administrative processes can enable extraordinary harm remains influential in contemporary governance scholarship.

Why Arendt's framework is invoked here:

This submission does not suggest contemporary Australia resembles Nazi Germany or totalitarian states. Rather, Arendt's analysis of institutional mechanisms provides a lens for examining how administrative systems can develop harmful characteristics even within democratic contexts when certain design features coincide.

Her focus on bureaucratic processes, diffusion of responsibility, and the gap between individual intent and systemic outcomes offers insights relevant to evaluating contemporary policy frameworks.

Core Insights from Arendt's Analysis

1. Persecution at Scale Requires Administrative Infrastructure

A central observation in Arendt's work is that **individual prejudice, however widespread, cannot produce systematic persecution without institutional machinery to implement it**.

Key insight:

> Individual antisemites have existed across centuries and societies. However, genocide or systematic persecution occurred only where states built administrative infrastructure capable of implementing policy at scale.

Implications for policy analysis:

When evaluating responses to antisemitism, this framework suggests examining not only the sincerity of protective intent but also:

- What administrative powers are being created?
- How are those powers constrained or reviewed?
- What happens if future governments reinterpret those powers?
- Does the machinery persist beyond the circumstances that prompted its creation?

Contemporary relevance:

Offshore detention (established 2001, persisting across multiple governments) and Robodebt (2016–2019, affecting 443,000 people despite internal warnings) both demonstrate how administrative systems can operate at scale with limited accountability once institutional frameworks are established.

---

2. The "Banality of Evil" and Bureaucratic Diffusion

Arendt's *Eichmann in Jerusalem* introduced the concept of the "banality of evil"—the observation that systematic harm can result from ordinary bureaucratic processes rather than extraordinary malice.

Eichmann's defence:

Adolf Eichmann, a senior Nazi official responsible for organising deportations during the Holocaust, claimed at trial that:

- He was following orders
- He was performing administrative functions
- He bore no personal responsibility for outcomes
- He never personally harmed anyone

Arendt's analysis was not that Eichmann lacked moral responsibility. Rather, she observed that **bureaucratic systems can diffuse responsibility such that individuals implementing harmful policies do not perceive themselves as making morally consequential decisions**.

The bureaucratic fragmentation of responsibility:

In complex administrative systems:

- **Person A** designs an algorithm
- **Person B** implements the system
- **Person C** maintains the database
- **Person D** sends automated notifications
- **Person E** handles appeals

No single person views themselves as causing harm. Each sees their role as technical, administrative, or procedural.

Relevance to automated decision-making:

The Robodebt Royal Commission documented precisely this dynamic. Officials described themselves as:

- implementing government policy,
- following ministerial direction,
- performing technical functions,
- executing administrative processes.

Yet collectively, the system produced 443,000 unlawful debt notices and caused documented harm including deaths.

Commissioner Catherine Holmes found that while individual officials may not have intended harm, the system's design enabled widespread unlawful outcomes—a finding consistent with Arendt's analysis of how bureaucratic diffusion operates.

---

3. Rightlessness and Legal Limbo

A third insight from Arendt's work concerns what she termed "rightlessness"—the condition of being outside legal protections.

The concept:

Arendt observed that stateless refugees between the World Wars existed in a legal void:

- No country recognised their citizenship
- No authority guaranteed their rights
- International law existed but proved unenforceable
- They depended entirely on administrative discretion

Legal limbo in contemporary Australian policy:

Offshore detention:

- Detainees neither fully subject to Australian law nor protected by host country law
- Limited access to courts
- Dependent on administrative decisions with minimal review

Robodebt:

- Recipients subject to unlawful debts but presumed valid until proven otherwise
- Burden on individuals to disprove averaged income
- System operated for years despite internal legal warnings

Schedule 2:

- Individuals may not be notified of adverse security assessments
- Limited avenues for review or appeal
- Assessments remain valid even if process was unlawful (statutory immunity)

The pattern across these cases is creation of situations where affected individuals have limited ability to invoke legal protections or challenge administrative decisions.

---

4. The Automation of Administrative Processes

Arendt observed that once bureaucratic systems are established, they tend to operate with institutional momentum independent of the circumstances that created them or the intentions of individual officials.

From *The Origins of Totalitarianism*:

> "Terror is the realization of the law of movement; its chief aim is to make it possible for the force of nature or of history to race freely through mankind, unhindered by any spontaneous human action."

Contemporary translation:

Once automated systems are deployed, they process cases regardless of:

- changing circumstances,
- individual variations,
- warning signals, or
- accumulating evidence of harmful outcomes.

Robodebt as illustration:

Despite:

- internal legal warnings,
- recipient complaints,
- media investigations,
- parliamentary questions, and
- advocacy organisation pressure,

the system continued operating for years. Momentum required Royal Commission intervention to stop.

Schedule 2 concern:

If algorithmic security assessments operate with statutory immunity and limited notification:

- How do affected individuals challenge errors?
- What triggers review of systemic problems?
- Who bears responsibility when automated processes produce harmful outcomes?

These questions mirror those raised by Robodebt, suggesting similar institutional vulnerabilities may exist.

Application to Contemporary Policy Design

Not Prediction—Pattern Recognition

This submission does **not** claim that:

- Schedule 2 will inevitably produce harms comparable to historical atrocities
- Officials designing the system intend harm
- Contemporary Australia resembles totalitarian regimes

Rather, it suggests that **certain institutional design patterns—when they appear together—create conditions where harmful outcomes become more likely, regardless of individual intentions**.

The concerning pattern includes:

1. **Automated or algorithmic decision-making** affecting individual rights
2. **Limited notification** of adverse decisions
3. **Barriers to effective review** or appeal
4. **Legal immunity** provisions protecting decisions from challenge
5. **Bureaucratic diffusion** where no individual feels responsible for outcomes

When these features converge, Arendt's framework suggests increased risk of systematic harm because:

- individuals affected cannot easily challenge decisions,
- officials implementing systems do not perceive themselves as making consequential moral choices,
- institutional momentum makes reversal difficult, and
- accountability mechanisms face structural obstacles.

Lessons from Robodebt

The Royal Commission's Findings

The Robodebt Royal Commission (2022–2023) provides a recent domestic case study relevant to evaluating institutional design.

Key findings:

1. **The scheme was unlawful from inception**
    - Income averaging had no legal basis
    - Senior officials were aware of serious legal doubts
    - Operation continued despite warnings
2. **Bureaucratic diffusion enabled continuation**
    - No single official felt personally responsible
    - Each person implemented "their part" of the system
    - Collective outcome was 443,000 unlawful debts
3. **Automated processes operated at scale**
    - Limited human oversight
    - Recipients faced burden of disproving averaged income
    - Volume of cases made effective review difficult
4. **Minimal accountability despite documented harm**
    - Royal Commission referred six individuals for potential prosecution
    - As of February 2026, no charges have been filed
    - Implicated officials faced limited consequences

Recommendation 17.1:

The Royal Commission specifically recommended that:

> "Automated decision-making systems affecting individuals' rights should include: clear review pathways, transparent explanations of how decisions are made, human oversight of automated processes, and accountability when systems produce harmful outcomes."

This recommendation directly addresses the institutional vulnerabilities identified in Arendt's framework: bureaucratic diffusion, automation without oversight, and limited accountability.

Schedule 2 Through an Institutional Design Lens

Structural Similarities to Robodebt

Schedule 2's design features include:

1. **Algorithmic or automated assessment processes** (sections 92E–92F)
2. **Removal of notification requirements** (section 36(1)(d) deleted)
3. **Limited review mechanisms** (no right to know basis of assessment)
4. **Statutory immunity from invalidity** (section 92G(3))

These features mirror the Robodebt architecture despite the Robodebt Royal Commission's explicit warning against such designs.

Timeline significance:

- Robodebt Royal Commission report: July 2023
- Schedule 2 passed Parliament: January 2026
- Time between: 28 months

The question this raises is whether Schedule 2's design reflects considered engagement with Robodebt's lessons or represents replication of the same structural vulnerabilities in a different policy context.

The Question of Protective Intent

Why Good Intentions Are Insufficient

A predictable response to this analysis is: "But Schedule 2 is designed to protect people, not harm them. The officials implementing it are well-intentioned. This is nothing like what Arendt described."

Arendt's framework specifically addresses this objection.

Her analysis was not about malicious intent. It was about how:

- systems designed for one purpose can be adapted to others,
- administrative processes can produce outcomes no individual intended,
- protective machinery can become harmful machinery when circumstances or leadership change, and
- bureaucratic structures persist beyond the conditions that created them.

Historical examples:

Many administrative systems that eventually enabled persecution were initially established for ostensibly protective or administrative purposes:

- Population registration (useful for census and service delivery; also used to identify targets)

- Movement controls (designed for security; also used to restrict disfavoured groups)
- Automated profiling (efficient administrative tool; also prone to bias amplification)

The lesson is not to avoid administrative systems. Liberal democracies require functioning bureaucracies.

The lesson is to design those systems with:

- robust oversight,
- transparent processes,
- effective review mechanisms,
- clear accountability, and
- recognition that future governments may use the same machinery differently.

Why This Matters for Social Cohesion

Trust in Institutions Depends on Perceived Fairness

The Commission's mandate includes examining threats to social cohesion. Institutional design affects social cohesion in several ways:

1. Perception of equal treatment

When administrative systems appear to operate with limited transparency or to affect certain communities disproportionately, broader populations may lose confidence that laws apply equally.

2. Trust in democratic processes

If policies persist despite evidence of harm, or if bipartisan support makes electoral accountability difficult, citizens may question whether democratic mechanisms function effectively.

3. Willingness to engage with institutions

Communities that feel subject to opaque or arbitrary decision-making may become alienated from civic participation and institutional engagement.

4. Potential for resentment

Policies perceived as privileging certain groups' security while subjecting others to enhanced scrutiny may produce inter-community tensions rather than cohesion.

These dynamics are not hypothetical. They have been documented across multiple policy contexts and contribute to the erosion of social trust that undermines cohesion.

Conclusion: Learning from Historical Patterns

This Part has examined Hannah Arendt's analytical framework not to claim that contemporary Australia faces totalitarian threats, but to consider what her scholarship suggests about institutional design.

Key insights:

1. **Persecution at scale requires machinery, not just hatred**
    - Protecting communities requires careful attention to what administrative powers are created
2. **Bureaucratic diffusion enables harm without malice**
    - Systems must include accountability mechanisms that function regardless of individual intent
3. **Rightlessness emerges when legal protections are removed**
    - Frameworks must preserve notification, review, and appeal rights
4. **Automation creates institutional momentum**
    - Safeguards must be built into design, not added after problems emerge

The Robodebt lesson:

Australia recently experienced how automated administrative systems operating with limited oversight can produce widespread harm despite officials' understanding that they were implementing lawful policy.

The Robodebt Royal Commission recommended specific safeguards to prevent recurrence.

The Schedule 2 question:

Does Schedule 2's design incorporate those safeguards, or does it replicate the architectural vulnerabilities that enabled Robodebt's harms?

This question is central to determining whether responses to antisemitism will strengthen democratic protections for all communities or create new institutional risks.

The Commission's opportunity:

Unlike Robodebt (which required a Royal Commission after harm occurred) or offshore detention (which has operated for 24 years without comprehensive inquiry), Schedule 2 can still be reformed before widespread impacts materialise.

Arendt's framework suggests that **the time to examine institutional design is before systems operate at scale, not afterward**.

This submission respectfully suggests that the Commission recommend strengthening Schedule 2's safeguards to align with:

- Robodebt Royal Commission Recommendation 17.1
- Established principles of administrative law
- Democratic accountability standards

- Recognition that protective machinery today may become harmful machinery tomorrow if not properly constrained

The goal is not to obstruct protection of Jewish Australians but to ensure that protection is achieved through institutional designs that preserve rather than undermine the democratic foundations on which all communities' safety ultimately depends.

PART VII: HISTORICAL PATTERNS IN AUSTRALIAN ADMINISTRATIVE POLICY

Introduction: Examining Recurring Structural Features

This Part examines whether certain administrative and legislative patterns observed across Australian policy domains over recent decades may assist the Commission in understanding contemporary challenges to democratic accountability and social cohesion.

The analysis focuses on three policy areas:

1. **Offshore detention arrangements** (2001–present)
2. **The Robodebt scheme** (2015–2019)
3. **Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026***

While these policies differ significantly in purpose, scope, and affected populations, they share certain structural features that may warrant consideration:

- creation of administrative frameworks where individuals face limited notification of adverse decisions;
- barriers to effective review or appeal;
- reliance on automated or algorithmically-assisted processes; and
- legislative design that appears to limit judicial oversight.

The submission does not assert that these similarities reflect deliberate replication or coordinated intent. Rather, it raises the question of whether recurring design patterns across otherwise distinct policy domains may indicate systemic governance challenges relevant to the Commission's mandate.

A note on framing: This analysis uses historical policy as a lens for examining institutional risk. It does not equate the humanitarian impact of offshore detention with other administrative systems, nor does it minimise the distinct harms experienced by affected individuals in each context. The focus is institutional architecture rather than moral equivalence.

I. The Tampa Incident and the Pacific Solution (2001)

Historical Context

On 26 August 2001, the Norwegian container vessel MV Tampa rescued 433 asylum seekers from a distressed boat in international waters. The vessel's captain altered course toward Christmas Island, the nearest Australian territory, in accordance with maritime safety obligations.

The Australian government denied the Tampa permission to enter territorial waters. On 29 August, Special Air Service Regiment personnel boarded the vessel. The Federal Court initially ruled that

the detention of rescued persons was unlawful, though this decision was later stayed pending appeal.

These events occurred during the lead-up to a federal election and contributed to the development of what became known as the Pacific Solution—a policy framework establishing offshore processing arrangements for asylum seekers who arrived by boat.

Political Dynamics and Narrative Construction

During the October 2001 election campaign, government officials made claims that asylum seekers aboard another vessel (SIEV-4) had thrown children overboard to compel rescue.

A subsequent Senate inquiry found that the claims were not supported by evidence. Photographs released during the campaign showed people in the water after their vessel had begun sinking, rather than parents deliberately endangering children.

Evidence presented to the inquiry indicated that senior officials were aware before the election that the initial reports were likely incorrect, yet public statements continued to reference the claims.

This episode illustrates how contested narratives about asylum seekers became embedded in public discourse during a period of heightened political attention to border security.

The Political Environment: One Nation and Electoral Strategy

The Tampa response occurred within a broader political context shaped by the emergence of Pauline Hanson's One Nation party in the late 1990s.

Hanson's 1996 maiden parliamentary speech included references to Australia being "in danger of being swamped by Asians" and questioned the sustainability of multiculturalism. One Nation polled between 8–10% nationally by 2001, creating competitive pressure in certain electorates.

Analysts have observed that the major parties responded to this pressure by adopting stricter rhetoric and policies regarding unauthorised maritime arrivals, while avoiding the explicitly racial language that characterised One Nation's discourse.

The term "boat people" became a shorthand category in political debate. Since most maritime arrivals originated from Middle Eastern, South Asian, or Southeast Asian countries, the category functioned as a proxy for ethnicity and national origin, even when framed in terms of arrival method rather than race.

Bipartisan Convergence

The Howard government won the November 2001 election. Labor initially opposed aspects of the Pacific Solution but gradually moved toward similar policy positions.

By 2012, the Gillard Labor government reopened offshore detention facilities on Nauru and Manus Island. In 2013, the Rudd government's "PNG Solution" declared that individuals arriving by boat would not be resettled in Australia—a position maintained by subsequent Coalition and Labor governments.

This bipartisan convergence is relevant to understanding why offshore detention has persisted across multiple changes of government. When both major parties support a policy framework, electoral mechanisms for policy change face structural obstacles.

Administrative Structure: The Pacific Solution Framework

The Pacific Solution, announced in September 2001, introduced several elements:

1. **Offshore processing centres** on Nauru and Manus Island (Papua New Guinea)
2. **Excision of certain Australian territories** from the migration zone
3. **Regional processing arrangements** without automatic resettlement in Australia
4. **Limited access to Australian courts** for individuals in offshore facilities
5. **Mandatory detention** with indeterminate duration

These arrangements created what legal scholars and human rights analysts have described as a condition of "legal ambiguity," where individuals were neither subject to full Australian law nor clearly protected by the domestic legal systems of host countries. International law frameworks proved difficult to enforce in practice.

Former officials and service providers have described offshore facilities as designed primarily for deterrence rather than timely processing. The intended message was that arriving by boat would not lead to Australian resettlement, regardless of refugee status.

II. Offshore Detention: Persistence Across Governments (2001–2025)

Continuity Despite Political Change

Since 2001, Australia has experienced multiple changes of government and prime minister. The offshore detention framework, however, has remained substantially intact.

Key features that have persisted across administrations include:

- **Maintenance of offshore processing centres** (with brief closures and reopenings)
- **Exclusion of boat arrivals from standard onshore processing**
- **Bipartisan political support** for the principle of offshore detention
- **Limited public reporting** on conditions and outcomes

This continuity raises questions about whether certain policy frameworks, once established, develop institutional and political momentum that makes reversal difficult regardless of which party governs.

Financial and Human Costs

Public reporting and parliamentary inquiries have documented substantial financial costs associated with offshore detention:

- **Estimated total expenditure**: Over $10 billion across 24 years

- **Per-person annual cost**: Estimates ranging from $573,000 to over $1 million per person detained offshore, compared to approximately $10,000–$15,000 per person for community-based onshore processing

Human costs have also been extensively documented:

- **Deaths in detention or shortly after release**: At least 14 individuals, including deaths classified as suicides
- **Mental health impacts**: Widespread reports of depression, PTSD, and self-harm
- **Family separation**: Prolonged uncertainty affecting children's development and family cohesion
- **Medical evacuation delays**: Cases where urgent medical treatment was delayed due to administrative or political obstacles

These outcomes have been examined by multiple inquiries, UN bodies, and human rights organisations. The Australian Human Rights Commission, parliamentary committees, and medical professionals have raised concerns about conditions in offshore facilities.

Corruption and Procurement Concerns

Offshore detention has also raised governance questions regarding contractor oversight and procurement practices.

The Paladin contract (2017–2019):

In February 2017, the Department of Home Affairs awarded Paladin Solutions Pty Ltd a $423 million contract (later expanded to $532 million over 22 months) to provide garrison and welfare services on Manus Island.

Parliamentary scrutiny revealed several concerns:

- Paladin had minimal prior experience in refugee services
- The company's registered office was a beach shack on Kangaroo Island
- The contract was awarded without competitive tender
- Questions arose about the company's capacity to deliver services at the contracted level

Senate inquiries sought documents and explanations but encountered significant delays and claims of commercial confidentiality.

The Canstruct allegations (2021):

In 2021, allegations emerged regarding Canstruct International, another contractor providing services on Nauru. These included:

- A director charged with immigration fraud
- Alleged provision of false documentation
- Questions about vetting and oversight processes

These cases illustrate governance challenges when large public expenditures occur in contexts with limited transparency and parliamentary oversight.

International Legal and Diplomatic Dimensions

Australia's offshore detention arrangements have attracted attention from international human rights bodies and foreign governments.

International Criminal Court consideration (2020):

In February 2020, the ICC Prosecutor's office concluded a preliminary examination into whether Australia's offshore detention policies constituted crimes against humanity, specifically:

- Deportation or forcible transfer of population
- Imprisonment or severe deprivation of liberty
- Torture, inhuman treatment
- Persecution

The Prosecutor declined to open a formal investigation, citing criteria related to gravity and complementarity. However, the examination itself indicated that offshore detention raised questions under international criminal law frameworks.

Diplomatic responses:

Papua New Guinea closed the Manus Island facility in 2017 following a Supreme Court ruling that detention there was unconstitutional under PNG law. This created operational challenges and required renegotiation of arrangements.

New Zealand repeatedly offered to resettle refugees from offshore detention—an offer the Australian government declined for many years, citing concerns about maintaining deterrence.

The Question of Ideological Commitment

The persistence of offshore detention despite high costs, documented harms, corruption concerns, and international criticism raises analytical questions about policy drivers.

Standard cost-benefit frameworks would suggest reconsidering policies that impose substantial financial and humanitarian costs. Yet offshore detention has continued across multiple governments regardless of expense or outcome.

This pattern suggests that the policy may serve purposes beyond administrative efficiency—functions related to electoral positioning, symbolic deterrence, or the maintenance of certain narratives about border security and national sovereignty.

Whether this constitutes "ideological commitment" (as some critics argue) or "consistent policy judgment" (as governments have maintained) is an interpretive question. What is observable is the policy's remarkable resistance to change despite evidence that would normally trigger reconsideration.

III. Robodebt: Automated Decision-Making Without Safeguards (2015–2019)

The Income Averaging Mechanism

Between 2016 and 2019, the Department of Human Services implemented an automated debt-raising system that compared welfare recipients' declared fortnightly income with annual income data from the Australian Taxation Office.

Where a discrepancy appeared, the system calculated an average fortnightly income across the year and generated a debt notice. This "income averaging" method assumed that income was distributed evenly across all fortnights, which was often factually incorrect for casual or seasonal workers.

The scheme affected approximately 443,000 individuals and raised debts totalling $1.8 billion.

Procedural Concerns

The Robodebt Royal Commission (2022–2023) found that the scheme incorporated several features that compromised procedural fairness:

1. **Reversal of onus**: Recipients were required to disprove debts rather than the Department proving they existed
2. **Limited notification**: Automated letters provided minimal explanation of how debts were calculated
3. **Barriers to review**: The volume of cases and complexity of disputing averaged income created practical obstacles to appeal
4. **Assumption of lawfulness**: Internal legal advice warning that the scheme was likely unlawful was not disclosed to Parliament or the public

Documented Harms

The Royal Commission documented serious consequences:

- **Deaths**: Evidence linking Robodebt to suicides, though causation is complex
- **Financial hardship**: Debts causing homelessness, inability to pay for food or medication
- **Mental health impacts**: Depression, anxiety, stress-related conditions
- **Social harms**: Damaged relationships, stigma, loss of trust in institutions

Commissioner Catherine Holmes found that the scheme was unlawful from inception and that senior officials were aware of serious legal doubts but proceeded nonetheless.

Accountability Outcomes

The Royal Commission referred six individuals for consideration of criminal charges, finding evidence of potential:

- Misconduct in public office
- Giving false or misleading evidence
- Perverting the course of justice

As of February 2026, no prosecutions have been initiated. The National Anti-Corruption Commission (NACC) declined to investigate the Robodebt referrals, a decision that itself became subject to review when the NACC Inspector found that the Commissioner had engaged in "officer misconduct" due to conflicts of interest during the decision-making process.

Recommendation 17.1: A Warning Against Automation Without Safeguards

The Royal Commission's Recommendation 17.1 stated:

> "The use of automated decision-making in the social welfare system should be subject to strict safeguards, including independent review, transparency about algorithms and decision-making processes, and clear accountability when automated systems produce harmful outcomes."

This recommendation reflected a clear finding: **automated systems affecting individuals' rights and entitlements require robust human oversight, transparent processes, and effective review mechanisms.**

IV. Schedule 2: Replication 28 Months Later (2023–2026)

The Legislative Framework

Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* amended the ASIO Act 1979 to introduce three significant changes:

1. **Removal of notification requirements** (section 36(1)(d) deleted): ASIO is no longer required to notify individuals of adverse security assessments in all cases
2. **Authorization of automated systems**: Sections 92E–92F permit the use of "approved software" or algorithmic processes in making security assessments
3. **Statutory immunity from invalidity** (section 92G(3)): Security assessments remain legally valid even if the decision-making process was unlawful, provided the person "reasonably believed" they were complying with the law

Similarities to Robodebt Architecture

The structural parallels between Robodebt and Schedule 2 include:

Feature	Robodebt	Schedule 2
Automated decision-making	Income averaging algorithm	Approved software for security assessment
Limited notification	Minimal explanation in debt letters	Notification requirement removed
Barriers to review	Burden on recipient to disprove debt	No right to know basis of assessment
Immunity provisions	Assumption assessments were valid	Statutory immunity from invalidity

Feature	Robodebt	Schedule 2
Timing	Implemented 2016–2019	Passed January 2026

The Robodebt Royal Commission reported in July 2023. Schedule 2 passed Parliament 28 months later, in January 2026.

This trajectory raises a question about institutional sustainability: are tools being created that will be available to future governments with different priorities? When far-right elements gain governmental influence—whether through parliamentary representation, coalition arrangements, or major party positioning—the algorithmic assessment capabilities created by Schedule 2 will be available to them. This is not speculation about current architects' intentions, but analysis of how institutional tools operate across political cycles.

Parliamentary Awareness and Concealment

Analysis of Hansard records indicates limited parliamentary awareness of Schedule 2's implications:

Senate debate (20 January 2026):

- Total speaking time on the entire Bill: Approximately 35 minutes
- Senators specifically discussing Schedule 2: Zero
- References to Robodebt Royal Commission: Zero
- References to Recommendation 17.1: Zero

House of Representatives debate (21 January 2026):

- Total speaking time: Approximately 40 minutes
- Members discussing Schedule 2 specifically: Zero
- Questions about algorithmic decision-making safeguards: Zero

The Explanatory Memorandum accompanying the Bill provided limited detail on Schedule 2's operation and did not reference the Robodebt Royal Commission findings or Recommendation 17.1, despite the clear relevance.

Inspector-General of Intelligence and Security Warning

The Inspector-General of Intelligence and Security (IGIS) issued a public statement in January 2026 (Submission 145 to parliamentary committee) expressing concern about Schedule 2's transparency and oversight provisions.

The IGIS noted that:

- Automated systems raise novel accountability questions
- Public confidence requires visibility of how algorithms operate
- Independent oversight mechanisms should be strengthened rather than weakened

This warning was issued **before** Parliament passed the legislation. The warning appears to have had no discernible effect on the legislative outcome.

V. Pattern Analysis: Recurring Structural Features

Common Design Elements Across Cases

Examining offshore detention, Robodebt, and Schedule 2 together reveals shared characteristics:

1. **Legal ambiguity or limbo**
    - Offshore: Neither fully subject to Australian law nor protected by host country law
    - Robodebt: Unlawful from inception but presumed valid until challenged
    - Schedule 2: Assessments valid even if process was unlawful
2. **Limited notification**
    - Offshore: Detainees often unaware of their legal status or pathways
    - Robodebt: Automated letters with minimal explanation
    - Schedule 2: Notification requirement removed
3. **Barriers to effective review**
    - Offshore: Limited court access, practical obstacles to legal representation
    - Robodebt: Burden on individuals to disprove averaged income
    - Schedule 2: No right to know basis of assessment
4. **Immunity from accountability**
    - Offshore: Political bipartisanship prevents inquiry
    - Robodebt: Senior officials faced no consequences despite Royal Commission referrals
    - Schedule 2: Statutory immunity even for unlawful processes
5. **Bipartisan support**
    - Offshore: Maintained across Liberal and Labor governments
    - Robodebt: Initially bipartisan, later disavowed
    - Schedule 2: Passed with support from both major parties

The Question of Causation

This submission does not claim that offshore detention directly caused Robodebt, or that Robodebt directly caused Schedule 2.

The question raised is different: **whether similar structural features appearing across distinct policy domains over 27 years reflects:**

- direct institutional learning and replication,
- parallel bureaucratic logic responding to similar political pressures,
- systemic governance weaknesses that recur across contexts, or
- some combination of these factors.

Answering this question definitively would require access to internal government communications and decision-making processes. What can be documented is the pattern itself: the repeated appearance of legal ambiguity, limited notification, barriers to review, and immunity provisions across otherwise distinct policy areas.

Arendt's Framework: Techniques and Rightlessness

Hannah Arendt's analysis in *The Origins of Totalitarianism* (1951) examined how democratic systems can develop authoritarian characteristics through administrative processes rather than explicit ideological transformation.

Key insights relevant to this analysis include:

1. **Rightlessness**: When legal systems create categories of people outside normal protections, those individuals become dependent on administrative discretion rather than enforceable rights.
2. **Bureaucratic diffusion**: Harmful outcomes can result from systems where no single actor feels responsible, as each person implements one part of a broader process.
3. **Normalisation through repetition**: Techniques that initially appear exceptional may become routine through repeated application across different contexts.

Arendt's framework does not predict inevitability. Rather, it identifies patterns that may indicate institutional vulnerabilities requiring attention.

The relevance here is not to compare contemporary Australia to totalitarian regimes, but to consider whether certain administrative patterns—legal limbo, reduced notification, limited review, immunity from accountability—may create conditions where harmful outcomes become more likely, regardless of individual officials' intentions.

VI. Social Cohesion Implications

Trust and Institutional Legitimacy

The Commission's mandate includes examining threats to social cohesion. The patterns documented in this Part raise several questions relevant to that mandate:

Perceived fairness:
When administrative systems appear to operate with limited transparency or accountability, affected communities may lose confidence in institutional fairness. This erosion of trust can extend beyond directly affected individuals to broader community perceptions.

Equal treatment:
If certain policies appear to target particular communities (offshore detention affecting predominantly non-European asylum seekers; Robodebt disproportionately affecting economically vulnerable citizens), questions arise about whether the rule of law applies equally to all.

Democratic responsiveness:
When policies persist despite documented harms, financial costs, and expert warnings—and when bipartisan support makes electoral accountability difficult—citizens may question whether democratic processes function effectively.

These concerns connect directly to social cohesion: communities that feel subject to arbitrary power, or that perceive differential treatment, may become alienated from broader society and civic participation.

The Deterrence Function and Community Division

Offshore detention was explicitly designed for deterrence—to send a message that arriving by boat would not result in resettlement. This deterrence function relied on visible punishment creating a cautionary example.

Some analysts have noted that visible punishment of one group can send signals beyond the immediate target population. When governments demonstrate capacity and willingness to maintain harsh policies despite costs and criticisms, it may communicate broader messages about:

- which communities' welfare matters to government,
- which legal protections are considered negotiable, and
- how democratic constraints operate in practice.

Whether these signals are intended or unintended, they may contribute to community perceptions that shape social cohesion.

VII. Why Royal Commission Into Offshore Detention May Be Necessary

The Accountability Deficit

Unlike Robodebt, which eventually received Royal Commission scrutiny after massive harm had occurred, offshore detention has operated for 24 years without comprehensive independent inquiry.

Parliamentary inquiries have examined specific aspects. Courts have ruled on individual cases. Human rights bodies have issued reports. But no Australian institution has conducted the type of systematic investigation that the Robodebt Royal Commission provided: compulsory evidence-gathering, public hearings, examination of internal documents, and findings about individual responsibility.

The reasons for this absence appear to relate to bipartisan political support. When both major parties maintain a policy, parliamentary mechanisms for investigation face obstacles. Opposition parties are unlikely to call for inquiries into policies they also support or have supported when in government.

This creates a democratic accountability gap: policies can continue indefinitely without the level of scrutiny typically applied to contested government actions.

What a Royal Commission Could Examine

A Royal Commission into offshore detention could investigate:

1. **Origins and decision-making**: How the Pacific Solution was developed, what advice was provided, how political and policy considerations interacted

2. **Human costs**: Comprehensive documentation of health outcomes, family impacts, and long-term effects on individuals
3. **Financial accountability**: Where $10+ billion was spent, whether procurement processes were adequate, what governance failures occurred
4. **Corruption and contractor oversight**: Full examination of Paladin, Canstruct, and other contracts with powers to compel evidence
5. **International law compliance**: Whether Australia met its obligations under refugee conventions, maritime law, and international criminal law
6. **Institutional learning**: Whether lessons from offshore detention influenced subsequent policies (Robodebt, Schedule 2)
7. **Individual responsibility**: Whether officials who made false claims (Children Overboard) or presided over harmful conditions should face accountability

Acknowledging Political Unlikelihood

This submission recognises that a Royal Commission into offshore detention is politically unlikely given bipartisan support for the policy.

However, recommending such an inquiry serves several purposes:

1. **Documentary record**: Creating a formal recommendation establishes that the question was raised and considered
2. **Moral clarity**: Acknowledging that some policies should face scrutiny even when political will is absent
3. **Precedent for future**: Demonstrating that Royal Commissions can examine bipartisan failures when circumstances eventually permit
4. **Accountability pressure**: Public recommendation may increase political costs of continued concealment

The Robodebt Royal Commission was also considered unlikely until political circumstances shifted. Recommendations made during periods when action seems impossible can become relevant when conditions change.

VIII. Connection to Schedule 2 and Current Legislative Response

The Pattern's Relevance to Antisemitism Policy

This Part's examination of offshore detention and Robodebt is not tangential to the Commission's mandate. Rather, it demonstrates:

How well-intentioned responses to genuine threats can create systemic risks when design features include legal ambiguity, reduced notification, limited review, and immunity from accountability.

How bipartisan support can prevent effective scrutiny of policies that might otherwise face parliamentary or judicial oversight.

How patterns established in one context can appear in subsequent, seemingly unrelated policy areas, suggesting systemic governance vulnerabilities rather than isolated failures.

The Antisemitism Bill's Schedule 2 provisions replicate design features that have proven problematic in earlier contexts. The question for the Commission is therefore not whether antisemitism requires response (it clearly does), but whether Schedule 2's architecture is likely to achieve protective goals without creating new vulnerabilities.

Breaking the Pattern

If Schedule 2 is allowed to operate without the safeguards recommended by the Robodebt Royal Commission, it may:

- produce harmful outcomes that only become visible after years of operation,
- prove resistant to reform due to bipartisan political dynamics,
- establish precedents for further expansion of automated decision-making without oversight, and
- contribute to erosion of trust in institutions among affected communities.

Alternatively, if the Commission recommends strengthening safeguards now—before harms accumulate—it may interrupt a 27-year pattern and demonstrate that democratic accountability can function effectively even in areas of bipartisan policy convergence.

Conclusion: The Commission's Opportunity

This Part has documented a pattern across three policy areas spanning 27 years: offshore detention, Robodebt, and Schedule 2.

The pattern includes:

- creation of legal limbo or ambiguity,
- limited notification of adverse decisions,
- barriers to effective review,
- immunity from accountability, and
- bipartisan political support preventing scrutiny.

In each case, substantial harms occurred or may occur. In each case, opportunities for earlier intervention existed but were not taken.

The Commission's examination of Schedule 2 represents a rare opportunity: the chance to prevent harm rather than document it retrospectively.

Unlike offshore detention (operating for 24 years) and Robodebt (operating for four years before Royal Commission), Schedule 2 can still be reformed before large-scale impacts occur.

The question is whether democratic institutions can learn from documented patterns and implement safeguards proactively, or whether the cycle of harm-then-inquiry will continue.

This submission respectfully suggests that the Commission recommend:

1. **Strengthening Schedule 2 safeguards** to align with Robodebt Royal Commission Recommendation 17.1
2. **Royal Commission into offshore detention** to address the 24-year accountability deficit (while acknowledging political unlikelihood)
3. **Systematic review of administrative frameworks** across government to identify and reform other systems exhibiting similar structural vulnerabilities

The Commission's findings on these matters could represent a turning point in Australian democratic governance—the moment when pattern recognition led to preventative action rather than retrospective regret.

PART VIII: SCALE AND INSTITUTIONAL RISK

The Fundamental Question for Liberal Democracy

Individual Violence and Institutional Power

On 14 December 2025, attackers opened fire at a Hanukkah celebration on Bondi Beach, killing 15 people. Police responded within minutes and stopped the attack shortly after it began. The perpetrators were neutralised and further harm prevented.

This horrific event prompted the introduction of the *Combatting Antisemitism, Hate and Extremism Bill 2026*, released publicly on 13 January 2026 and passed by both Houses of Parliament on 20–21 January 2026.

The stated objective of the legislation is to protect Jewish Australians and respond to rising antisemitic violence. That objective is legitimate and important. Jewish Australians, like all Australians, are entitled to safety, dignity, and protection under law.

This submission does not question the seriousness of antisemitism. Rather, it asks a structural governance question:

How should liberal democratic systems evaluate risk when responding to individual acts of violence?

Specifically:

- What risks arise from individual or group violence?
- What risks arise from institutional machinery created by the state to respond to that violence?

Traditional liberal democratic theory recognises that these risks are different in nature. Individual violence can cause immediate and profound harm. Institutional power, however, operates at a broader scale and persists beyond the circumstances that gave rise to it. When poorly constrained, state systems can produce systemic harm across entire populations.

This Part therefore examines **scale, proportionality, and institutional risk**.

Scale and the Asymmetry of Power

The Bondi attack illustrates the devastating impact that individual violence can have. Fifteen people lost their lives in a short period before existing security responses intervened. Such attacks require prevention, investigation, and accountability.

However, liberal democratic systems have historically developed safeguards not only to address individual threats but also to constrain the potential reach of state power. The rationale is straightforward:

- Individual actors can cause harm within limited operational capacity.
- State institutions possess legal authority, administrative infrastructure, and technological capability capable of affecting thousands or millions of people simultaneously.

This distinction does not minimise individual violence. Rather, it reflects a structural reality recognised across democratic theory: **the scale of institutional harm may exceed the scale of individual harm when safeguards fail.**

Modern constitutional systems therefore emphasise:

- separation of powers;
- procedural fairness;
- transparency;
- independent oversight; and
- judicial review.

These safeguards exist not because governments are inherently malicious, but because institutional power endures beyond any single government or political moment.

Lessons from Recent Australian Experience

Australia’s recent experience with the Robodebt scheme provides a relevant domestic example of institutional scale.

Between 2016 and 2019:

- automated processes generated hundreds of thousands of debt notices;
- legal safeguards failed;
- meaningful human review was limited; and
- systemic harm occurred before corrective action was taken.

The Robodebt Royal Commission documented how administrative automation, combined with insufficient oversight, enabled large-scale impacts on individuals despite internal warnings.

Recommendation 17.1 addressed precisely this risk: automated systems affecting individuals should include clear review pathways, transparent explanations, and mechanisms for independent scrutiny.

The relevance here is institutional, not ideological. Robodebt demonstrates that:

> **systems designed with legitimate policy objectives can nonetheless produce widespread harm when structural safeguards are absent or weakened.**

Evaluating Legislative Risk in Context

The Antisemitism legislation aims to address real and serious concerns. However, when evaluated through a liberal-democratic risk lens, the design of institutional powers becomes as important as the problem they seek to solve.

Key considerations include:

- automated intelligence assessment processes;
- executive decision-making authority;
- limited notification pathways; and
- reduced opportunities for independent review.

The question for this Commission is therefore not whether antisemitism should be addressed — it clearly must — but whether the institutional architecture chosen introduces risks disproportionate to the problem.

A governance framework must weigh:

1. **Immediate security benefits**, and
2. **Long-term institutional resilience.**

History demonstrates that institutional mechanisms created during moments of crisis may persist long after the initial threat has changed.

Liberal Democratic Risk Management

Liberal democratic systems often accept a difficult balance:

- Some level of individual risk may remain despite robust security measures.
- However, institutional safeguards are preserved to prevent systemic or arbitrary exercises of state power.

This balance is not an expression of indifference toward victims. It reflects recognition that:

> **institutional overreach, once normalised, can reshape democratic governance in ways that are difficult to reverse.**

Accordingly, democratic design prioritises:

- maintaining human decision-making;
- ensuring procedural fairness;
- preserving judicial oversight; and
- limiting executive discretion where rights may be affected.

Why Institutional Design Matters

A recurring pattern identified in both international history and Australian administrative practice is that **institutional systems develop operational momentum independent of their original intent**.

Policies introduced with protective or preventative aims may evolve over time through:

- administrative expansion;
- political pressure;
- technological scaling; or
- shifting security priorities.

The Robodebt Royal Commission documented how institutional processes — rather than malicious intent — enabled widespread harm. Officials often described themselves as implementing systems rather than exercising discretionary judgment.

This Part therefore approaches the Antisemitism legislation through a design-based lens:

> **the question is not motive, but architecture.**

Institutional Risk and Social Cohesion

Another dimension relevant to this Commission is social cohesion.

Policies perceived as disproportionately targeting particular communities may unintentionally increase social division, even when introduced for protective purposes. Community trust in institutions is a key element of effective counter-extremism strategy.

Where automated or opaque processes operate without clear review pathways, individuals may feel subject to decisions they cannot understand or challenge. Over time, this may undermine confidence in democratic institutions — an outcome inconsistent with the Bill’s stated objectives.

The Commission’s Role

This Commission is uniquely positioned to consider not only immediate harms but also institutional design.

The issue presented here is not whether government should respond to antisemitism — it must. Rather, it is whether the legislative response:

- aligns with recent lessons from Robodebt;
- incorporates sufficient safeguards;
- maintains democratic accountability; and
- avoids creating long-term structural risks.

The Inspector-General of Intelligence and Security has already encouraged consideration of public visibility and oversight in automated arrangements. That perspective reinforces the importance of examining institutional architecture carefully.

Conclusion: A Liberal-Democratic Principle

Liberal democracy rests on a principle developed through historical experience:

> **effective security requires both protection from violence and protection from the unintended consequences of concentrated state power.**

This submission therefore does not argue against security measures. Instead, it suggests that institutional responses must be designed in a manner that preserves:

- proportionality;
- transparency;
- independent oversight; and
- democratic legitimacy.

Subsequent Parts of this submission analyse the legislative framework in detail and propose alternative governance models demonstrating that security objectives and robust safeguards are compatible.

PART IX: PARLIAMENTARY CONSIDERATION AND LEGISLATIVE TRANSPARENCY

Introduction: Examining the Information Available to Parliament

On 20–21 January 2026, the Australian Parliament debated and passed the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*. Schedule 2 of that legislation introduced amendments to the Australian Security Intelligence Organisation Act 1979 (ASIO Act), establishing a framework for algorithmic assessment processes affecting security determinations.

This Part examines:

- The information provided to Parliament through the Explanatory Memorandum
- The substance of parliamentary debate as recorded in Hansard
- Constitutional analysis by independent experts
- Whether the legislative process met transparency standards appropriate for frameworks affecting individual rights

The submission does not allege improper intent by drafting officials or parliamentarians. Rather, it examines whether the transparency and disclosure practices surrounding Schedule 2 met standards appropriate for legislation introducing novel administrative frameworks with significant implications for individual rights and democratic accountability.

I. The Robodebt Context and Selective Application in the Explanatory Memorandum

The Robodebt Royal Commission Findings (July 2023)

Twenty-eight months before the Antisemitism Bill was introduced, Commissioner Catherine Holmes handed down findings from the Royal Commission into the Robodebt Scheme.

Recommendation 17.1 stated:

> "The Commonwealth should not pursue or continue the development of automated decision-making systems where there are known or likely deficiencies in the data or the systems are otherwise unlikely to comply with relevant legal requirements."

The Royal Commission documented how automated systems operating with:

- limited human oversight,

- inadequate notification of affected persons,
- immunity provisions limiting judicial review, and
- "reasonable steps" requirements that proved difficult to enforce

produced widespread unlawful outcomes despite officials' belief they were implementing lawful policy.

The Asymmetry in the Explanatory Memorandum

The Explanatory Memorandum (EM) demonstrates awareness of Robodebt Royal Commission findings—but applies them selectively.

For ACIC automation (criminal intelligence assessments for firearms background checks), paragraph 287 states:

> "The clauses in this Division are based on the Government's response to recommendation 17.1 of the Royal Commission into the Robodebt Scheme Report and the Commonwealth Ombudsman's Automated Decision-making Better Practice Guide. They are intended to ensure automated systems comply with administrative law principles of legality, fairness, rationality and transparency."

This explicit acknowledgment appears in Division 3 of the ACC Act amendments, which establish automated criminal intelligence assessments for firearms licensing purposes.

For ASIO automation (sections 92E–92G), paragraphs 354–362 contain:

- No reference to the Robodebt Royal Commission
- No reference to Recommendation 17.1
- No reference to the Commonwealth Ombudsman's Automated Decision-making Better Practice Guide
- No statement that the provisions are designed to ensure compliance with "administrative law principles of legality, fairness, rationality and transparency"

Instead, the EM characterises ASIO automation provisions primarily in terms of:

- "efficiencies for ASIO in appropriate circumstances" (paragraph 354)
- "high volume of assessments" (paragraph 356)
- "ongoing oversight of the Inspector General" (paragraph 358)

The Unexplained Differential Treatment

This creates an unexplained asymmetry in how Robodebt lessons were applied across the same legislation:

Framework Feature	ACIC Automation (ACC Act)	ASIO Automation (Sections 92E–92G)
Robodebt RC 17.1 Referenced	✓ Yes (paragraph 287)	✗ No
Ombudsman Guide Referenced	✓ Yes (paragraph 287)	✗ No
Explicit Administrative Law Principles	✓ "legality, fairness, rationality and transparency"	✗ Not stated

Framework Feature	ACIC Automation (ACC Act)	ASIO Automation (Sections 92E–92G)
Immunity from Invalidity	✗ Not included	√ Section 92G(3)
Notification Requirement Deleted	✗ Not applicable	√ Section 36(1)(d) deleted

Questions Raised by the Asymmetry

The differential application raises several transparency questions:

1. Why different standards for different agencies?

ACIC assessments inform firearms licensing decisions. ASIO assessments determine:

- visa cancellations and refusals
- citizenship denials
- security clearances
- passport cancellations
- access to sensitive information

If Robodebt safeguards are necessary for criminal intelligence assessments, why are they not explicitly applied to security assessments with potentially greater individual consequences?

2. What policy analysis supported the distinction?

The EM provides no explanation for why:

- ACIC automation required explicit reference to Robodebt lessons
- ASIO automation did not

3. Was Parliament informed of the differential treatment?

The structural presentation of the EM may have obscured this asymmetry:

- ACIC provisions appear in one schedule (Schedule 1, Division 3)
- ASIO provisions appear in a different schedule (Schedule 2, Item 30)
- No comparative analysis highlights the different approaches

What This Asymmetry Suggests

The selective application of Robodebt safeguards is **more concerning** than simple omission would be:

If Robodebt had been entirely forgotten:

- This might suggest oversight or time pressure
- Correctable through amendment

If Robodebt was remembered for ACIC but not ASIO:

- This suggests a deliberate choice about which frameworks require explicit safeguards
- This requires explanation for why ASIO automation is treated differently
- This raises questions about whether the same standards of administrative law compliance apply across Commonwealth agencies

The submission does not assert deliberate concealment. However, it suggests that transparency standards for legislation introducing automated decision-making frameworks should include:

- Consistent application of Royal Commission lessons across similar frameworks
- Explicit explanation when different agencies receive different safeguard standards
- Disclosure of policy analysis supporting differential treatment

---

II. Characterisation of the Automation Framework in the Explanatory Memorandum

How the Explanatory Memorandum Describes Schedule 2

The EM characterises the automation provisions using several recurring framings:

1. "Clarification" language (paragraphs 27, 33):

> "Further amendments would **clarify** the use of automation for specified assessment action."

Analysis: The term "clarify" suggests explaining existing authority rather than creating new powers. However, sections 92E–92G introduce statutory frameworks that did not previously exist in the ASIO Act.

2. "Beneficial to the applicant" framing (paragraph 27):

> "This automation would be beneficial to the applicant, as it would only be used where the system could not return positive results."

Analysis: This framing emphasises efficiency benefits while providing limited information about what happens when systems do return concerning results or operate incorrectly.

3. "High-volume efficiency" justification (paragraph 289):

> "These amendments would enable high-volume assessments to be processed efficiently while maintaining appropriate safeguards."

Analysis: The reference to "appropriate safeguards" is not followed by detailed explanation of what those safeguards are or how they operate in practice.

Features Not Prominently Disclosed

Several significant features of Schedule 2 receive limited or technical explanation in the EM:

1. **Statutory immunity provision** (section 92G(3)): Assessments remain valid even if the automated process was defective, provided officials took "reasonable steps"
2. **Third-party deployment authority** (section 92E(1)): ASIO may "arrange for the use" of approved software—potentially enabling private sector involvement

3. **Ministerial scope control** (section 92F(5)): Minister determines "specified assessment action" via non-legislative instrument—no parliamentary oversight of scope expansion
4. **Retrospective substitution** (sections 92E(6)–(8)): Automated assessments can retroactively replace human assessments
5. **Notification removal** (deletion of section 36(1)(d)): Individuals may not be informed they've been assessed
6. **"Reasonable steps" standard**: The primary safeguard relies on a standard the Robodebt Royal Commission found difficult to enforce

The "Firearms Licensing" Frame

Throughout the EM, Schedule 2's amendments to the ASIO Act are consistently framed in terms of firearms licensing:

- Paragraph 28: "firearms background check framework"
- Paragraph 268: "firearms licensing decisions"
- Paragraph 348: "in relation to firearms licensing"

However, the statutory language indicates broader potential application:

- Sections 92E–92G are inserted into the ASIO Act permanently
- They apply to "specified assessment action"—defined by ministerial instrument
- Section 92E(1) applies to ASIO's functions generally ("specified assessment action that may, or must, be taken by the Organisation under this Act")
- The EM notes automation may apply to "assessments under Parts IV and Part IVA" giving examples but not limiting scope

The gap between framing and statutory scope:

The EM's consistent firearms framing may have encouraged readers to understand Schedule 2 as a narrow, context-specific reform when the statutory language permits broader application.

This is not necessarily improper—legislation often begins with specific applications and expands over time. However, transparency standards might suggest that EMs should explicitly note when statutory language permits broader application than initial policy intent contemplates.

III. Parliamentary Debate: Analysis of Hansard Records

Overview of Debate

The Bill was debated over two days:

- **Senate**: 20 January 2026 (95 pages of Hansard)
- **House of Representatives**: 21 January 2026 (64 pages of Hansard)
- **Total**: 159 pages of parliamentary debate

Primary Focus of Debate

Analysis of Hansard records reveals debate concentrated on three main topics:

1. Firearms provisions:

- Gun buyback schemes and compensation
- Firearms licensing restrictions
- Sporting shooters' rights and rural firearm use
- Comparisons to John Howard's 1996 reforms

2. Antisemitism and community safety:

- Response to the Bondi Beach attack
- Protection of Jewish communities
- Naming of Islamic extremism
- Royal Commission timing

3. Process concerns:

- Rushed legislative timeline
- Limited consultation period (7 days)
- Volume of submissions to Parliamentary Joint Committee
- Bill splitting from omnibus to separate instruments
- Free speech implications

Statistical Analysis of Hansard Content

Word frequency analysis across 159 pages of combined Hansard:

Term	Frequency	Context
"Firearms" / "guns"	400+	Buyback, licensing, sporting shooters
"Antisemitism" / "Jewish"	300+	Attack, victims, community safety
"Islamic" / "extremism"	200+	Ideology, terrorism, hate preachers
"Automation"	**0**	—
"AI" / "artificial intelligence"	**1**	Senator McDonald only
"Robodebt"	**1**	Senator McDonald only
"Algorithm" / "algorithmic"	**0**	—
"Section 92E" / "92F" / "92G"	**0**	—
"Immunity from invalidity"	**0**	—
"Retrospective substitution"	**0**	—
"Computer programs"	**0**	—
"Non-legislative instrument"	**0**	—

Observation: Schedule 2's automated assessment framework received virtually no substantive parliamentary discussion despite its novel character and relationship to recent Royal Commission findings.

The Single Exception: Senator McDonald's Statement

Senator McDonald (Queensland, Nationals)
Senate Hansard, 20 January 2026, page 27:

> "This bill also allows us the use of AI to assist with decision-making on intelligence findings. That contravenes a recommendation of the robodebt royal commission, but it's in the legislation."

Context:

This statement represents the only substantive mention across 159 pages of Hansard connecting:

- AI-assisted decision-making
- Robodebt Royal Commission
- Recommendation 17.1
- Schedule 2 provisions

The comment appeared during a broader speech opposing the Bill on firearms grounds. It was not:

- Responded to by government
- Picked up by other speakers
- Debated substantively
- Addressed in committee reports

Significance:

The most consequential structural connection—that Schedule 2 raises concerns similar to those documented in the Robodebt Royal Commission—was identified by one senator and did not generate further parliamentary discussion.

What Members Appear to Have Understood They Were Enacting

Parliamentary statements suggest many members understood the Bill as:

1. Intelligence sharing for firearms licensing

Senator Watt (Government) (Senate Hansard, 20 January 2026, p. 15):

> "Intelligence held by the Australian Security Intelligence Organisation and the Australian Criminal Intelligence Commission would be shared securely and underpin firearms licensing decisions."

2. Background checks for firearm applicants

Prime Minister Albanese (House of Representatives Hansard, 20 January 2026, p. 2):

> "The bill will also lay the foundation for a new Commonwealth background checking framework to inform decisions to issue or renew firearm licences by the states and territories."

3. Temporary measures responding to specific threat

Member for Bendigo, Ms Chesters (House of Representatives Hansard, 20 January 2026, p. 3):

> "Following the devastating impact of the Bondi terrorist attack in Sydney, our government is rightfully reviewing the gun laws in Australia, as Australians demand."

Observations:

These statements characterise Schedule 2 as:

- Limited to firearms contexts
- Information-sharing rather than automated assessment
- Responsive to immediate crisis
- Temporary review of gun laws

The statutory language, however:

- Inserts provisions into ASIO Act permanently
- Authorises "specified assessment action" defined by ministerial instrument
- Applies to "specified assessment action that may, or must, be taken by the Organisation under this Act" (s 92E(1))
- Creates automated assessment authority with statutory immunity

The gap between stated understanding and statutory effect:

Whether this gap reflects:

- Members relying on EM framing without examining statutory detail
- Time constraints preventing detailed technical analysis
- Political focus on firearms and antisemitism overshadowing administrative machinery
- Or some combination of these factors

cannot be definitively determined from Hansard alone. However, the gap raises questions about whether the legislative process provided adequate opportunity for informed consideration of Schedule 2's implications.

The Commission may encounter pressure to accept Schedule 2 as enacted, on the basis that current political alignments ensure appropriate implementation. However, relying on political influence rather than institutional safeguards is precisely the vulnerability these provisions create. The question is not who controls the political narrative today, but what tools will exist when different people control government tomorrow.

IV. The Parliamentary Committee Process

Timeline and Constraints

The Parliamentary Joint Committee on Intelligence and Security (PJCIS) examined the original omnibus Bill under significant time constraints:

- **Exposure draft released**: Tuesday 14 January (9am)
- **Submissions due**: Thursday 16 January (5pm)
- **Hearings**: Friday 17 January

- **Report due**: Monday 20 January (morning)
- **Debate/passage**: Monday 20 January (afternoon)

Total consultation period: 7 days

UN Special Rapporteur's Concern

Member for Indi, Dr Haines, quoting the UN Special Rapporteur's submission:

> "It is the experience of the Special Rapporteur's mandate worldwide that hasty law-making is prone to lead to over-reach, unintended consequences, and violations of human rights law, which does not well serve the victims of terrorist attacks or protect society."

Committee Report Findings

The PJCIS report tabled on the morning of passage contained limited substantive analysis of Schedule 2's automated intelligence framework. The Committee's focus appeared directed primarily toward:

- Hate speech provisions
- Group banning mechanisms
- Migration Act amendments
- Firearms licensing coordination

Schedule 2's ASIO Act amendments received minimal detailed examination in the final report.

V. Last-Minute Constitutional Architecture Change: State Law Incorporation

The Bill Split and Substantive Amendment

On the morning of 20 January 2026, hours before final passage, the government split the omnibus Bill into two separate Acts. This division introduced a substantive change to the constitutional architecture that had not been examined during the inquiry process.

What Changed: Federal Incorporation of State Racial Hatred Laws

Original omnibus Bill structure:

"Hate crime" triggers for group banning defined exclusively by reference to Commonwealth offences.

Enacted structure (Criminal and Migration Laws Act 2026):

"Hate crime" includes conduct that would constitute an offence under **specified State or Territory racial hatred laws**.

Location: Schedule 1, Part 5.3B, Division 114A

Practical effect:

Commonwealth powers (group banning, criminal membership offences, visa cancellation) now depend on State and Territory law definitions that vary by jurisdiction.

Professor Anne Twomey's Constitutional Analysis

Professor Anne Twomey (University of Sydney), writing in *The Conversation* (21 January 2026):

> "It would also include conduct that would constitute a specified state or territory offence."
>
> "This reliance on state offences makes the law very messy."
>
> "No relevant laws are listed for Tasmania or the Northern Territory. This means that whether a group can be banned on this basis may depend on where the conduct took place."

On her YouTube channel (*Constitutional Clarion*):

> "The reference to the State laws was deliberate. It was a means of getting around the federal offence being dropped."

Constitutional Questions Raised

1. Geographic variation in federal law:

Federal criminal consequences now vary by jurisdiction:

- Conduct lawful in Tasmania may trigger federal group banning elsewhere
- Different State laws have different elements, defences, and thresholds
- Whether a group can be banned may depend on where alleged conduct occurred

Example variations:

- NSW: *Racial Discrimination Act 1977* s 20D (incitement to hatred)
- Victoria: *Racial and Religious Tolerance Act 2001* s 24 (incitement to serious contempt)
- Queensland: *Anti-Discrimination Act 1991* s 131A (incitement to hatred—tied to threatening physical harm)
- Tasmania: No equivalent law listed

Twomey's observation:

> "In the listed offences from Queensland, South Australia and the ACT, incitement to racial hatred is tied to threatening physical harm, whereas in New South Wales, Victoria and Western Australia, no threat of harm is required."

2. Ministerial determination replacing judicial process:

Twomey's core concern:

> "This leaves it up to the minister to decide what was done and by whom, whether they had the necessary intent, whether their conduct can be attributed to the group, whether

> any defences apply, and whether the conditions of the law of the relevant jurisdiction have been met.
>
> Ordinarily, we leave such assessments to independent courts and judges. For example, should a minister be the one deciding whether a defence of acting in good faith should apply, when the minister has a political interest in banning a particular group?"

Elements traditionally determined by courts now subject to ministerial assessment:

- *Actus reus*: What conduct occurred?
- *Mens rea*: What was the intent?
- Attribution: Can conduct be attributed to the group?
- Defences: Do statutory defences apply?
- Jurisdictional elements: Which State law applies?

3. Retrospective application concerns:

Twomey:

> "Conduct can be a 'hate crime' even though it happened in the past when it wasn't a crime. It is enough for the minister to be satisfied on reasonable grounds that the group has engaged in or been associated with the conduct constituting a 'hate crime'."

Conduct lawful when performed may retrospectively become the trigger for:

- Group banning determinations
- Criminal membership offences
- Visa cancellation
- Permanent exclusion from Australia

4. Impact on political communication:

Attorney-General Rowland's ABC 7.30 Interview:

When asked whether groups could be banned for criticising Israel if such criticism caused Jewish Australians to feel intimidated:

Rowland: "**If those criteria are satisfied, then that is the case**."

Senator Thorpe's amendment (rejected 43-12) (Senate Hansard, 20 January 2026, p. 94):

> "As per the judgement of the Federal Court in *Wertheim v Haddad* [2025] FCA 720, criticism of the practices, policies, and acts of the state of Israel, the Israeli Defence Forces or Zionism is not inherent criticism of Jewish people and is protected political speech, and not hate speech."

Twomey's assessment:

> "This leaves uncertain what conduct is intended to be caught. Freedom of political communication by those who wish to protest against the conduct of a nation's government could potentially be chilled."

Timing and Democratic Process Concerns

Key timeline facts:

- State law incorporation appeared in the split Bill on the **morning of 20 January**
- PJCIS inquiry examined Commonwealth-offence-only structure
- 7,000 submissions responded to Commonwealth-offence-only structure
- No opportunity for:
    - Constitutional analysis of federal-State architecture
    - Submissions on geographic variation implications
    - Expert evidence on federalism concerns
    - Public scrutiny of ministerial determination replacing judicial process

Observation:

Parliament enacted a constitutional structure substantively different from what was examined during the inquiry process. The change was introduced hours before passage without opportunity for independent analysis or public consideration.

VI. The Inspector-General's Warning

IGIS Submission to Parliamentary Committee

The Inspector-General of Intelligence and Security (IGIS) issued Submission 145 to the PJCIS on 15 January 2026, expressing concern about Schedule 2's automated assessment framework and recommending enhanced oversight mechanisms.

IGIS observations included:

- Automated systems raise novel accountability questions
- Public confidence depends on visibility of how algorithms operate
- Independent oversight mechanisms should be strengthened when automation is introduced
- Transparency requirements are particularly important for intelligence assessment contexts

IGIS recommendations:

- Mandatory consultation with IGIS before deploying automated intelligence systems
- Enhanced transparency about algorithmic categorisation processes
- Clear review mechanisms for affected persons

Parliamentary Response

Timeline:

- IGIS submission: 15 January 2026
- Bill passed: 20–21 January 2026
- Interval: 5–6 days

IGIS recommendations in enacted legislation:

The enacted framework did not incorporate the specific safeguards IGIS recommended regarding automated systems.

Observation:

An independent oversight body with specialist expertise in intelligence operations flagged concerns about Schedule 2's design and recommended specific safeguards. These recommendations were issued before Parliament passed the legislation but were not reflected in the final enacted framework.

This sequence raises questions about:

- Whether sufficient weight was given to specialist oversight body warnings
- What mechanisms exist to ensure expert technical concerns receive parliamentary attention
- Whether the compressed timeline prevented adequate consideration of IGIS recommendations

---

VII. International Comparisons: Transparency Standards for Automated Government Systems

European Union AI Act

The EU's regulatory framework for artificial intelligence in government includes:

- Mandatory transparency about algorithmic systems affecting individuals
- Requirements for human oversight of high-risk AI applications
- Explicit rights for affected persons to receive explanations
- Independent regulatory bodies with enforcement authority
- Public disclosure of impact assessments before deployment

United Kingdom Algorithmic Transparency Standard

The UK framework requires government departments using algorithms to:

- Publish descriptions of algorithmic systems
- Explain what data is used and how decisions are reached
- Provide information about testing and evaluation processes
- Establish clear accountability mechanisms
- Enable public and parliamentary scrutiny

Schedule 2 in Comparative Context

Australia's Schedule 2 framework differs from these international approaches by:

- Not requiring publication of algorithmic methodology
- Not mandating notification to affected persons
- Providing statutory immunity rather than enhanced review rights
- Relying on "reasonable steps" rather than specific transparency requirements

- Minimising rather than maximising public visibility into automated processes

Analytical question:

Whether Australia's approach reflects considered judgment that our constitutional and security context requires different safeguards, or whether Schedule 2 was developed without full engagement with international best practices in algorithmic governance that have emerged following concerns similar to those raised in the Robodebt Royal Commission.

VIII. Implications for Democratic Accountability

The Gap Between Statutory Effect and Parliamentary Understanding

The analysis above suggests a potential gap between:

- **What Schedule 2 authorises**: Broad algorithmic assessment framework with ministerial scope control, immunity provisions, and potential third-party deployment
- **How it was presented**: Firearms licensing efficiency measure with appropriate safeguards
- **Parliamentary awareness**: Minimal substantive debate on automated assessment provisions

This gap raises questions about whether Parliament had sufficient information to meaningfully assess:

- The relationship between Schedule 2 and Robodebt vulnerabilities
- The scope of ministerial authority to expand "specified assessment action"
- The practical effects of immunity provisions on review rights
- The implications of removing notification requirements
- The adequacy of "reasonable steps" as primary safeguard when the Robodebt Royal Commission found this standard difficult to enforce

Standards for Legislative Transparency

Democratic theory suggests that informed parliamentary consent requires:

1. **Accurate characterisation** of what legislation does in practice
2. **Disclosure of relevant context**, including recent findings on similar frameworks
3. **Plain-language explanation** of practical implications for affected persons
4. **Explicit acknowledgment** when statutory scope exceeds immediate policy intent
5. **Adequate time and information** for substantive parliamentary debate and expert review

Where legislation introduces novel administrative frameworks with significant implications for individual rights, these transparency standards become particularly important.

Comparative Assessment

The legislative process for Schedule 2 included:

- ✗ 7-day consultation period (international standards typically require months)
- ✗ 185-page EM covering multiple distinct reforms (complexity barrier)

- ✗ No reference to Robodebt Royal Commission despite direct relevance
- ✗ Consistent firearms framing despite broader statutory language
- ✗ Last-minute constitutional architecture change (State law incorporation)
- ✗ IGIS warnings issued but not incorporated before passage
- √ Parliamentary committee review (though constrained by timeline)
- √ Public submissions accepted (though processing time limited)

Question for consideration:

Whether this process met standards appropriate for legislation creating automated intelligence assessment frameworks with immunity provisions, particularly given Australia's recent experience with Robodebt demonstrating how automated administrative systems can produce widespread harm when safeguards prove inadequate.

IX. Connection to Social Cohesion

How Legislative Process Affects Community Trust

The Commission's mandate includes examining threats to social cohesion. Legislative process affects social cohesion through several pathways:

1. Perceived legitimacy of law

When legislation is perceived as having been passed without adequate scrutiny or transparency, affected communities may question its legitimacy regardless of substantive merit.

2. Trust in democratic institutions

If significant administrative powers are introduced with minimal parliamentary debate or public understanding, citizens may lose confidence that democratic processes function effectively to constrain executive authority.

3. Community preparedness to comply

Laws perceived as inadequately considered or explained may face greater resistance from affected populations than laws resulting from transparent, deliberative processes.

4. Basis for future accountability

When Parliament is not fully informed about legislation's implications, it becomes difficult to hold governments accountable for outcomes—because what was authorised versus what was understood may diverge significantly.

Specific Concerns for Affected Communities

If Schedule 2 is perceived as:

- having been passed with minimal scrutiny of its automation provisions,
- potentially affecting certain communities disproportionately through algorithmic assessment,
- operating with limited transparency or accountability mechanisms, and

- replicating structural features that proved problematic in Robodebt

this perception may contribute to feelings of alienation or mistrust among those communities—outcomes inconsistent with social cohesion objectives.

Conclusion: Questions for the Commission's Consideration

This Part has examined the information provided to Parliament about Schedule 2 and the parliamentary process that preceded its enactment.

Key observations:

1. **Robodebt safeguards applied selectively** — Explanatory Memorandum explicitly referenced Robodebt RC Recommendation 17.1 for ACIC automation but not for ASIO automation (sections 92E–92G), creating unexplained differential treatment across the same legislation
2. **Framing emphasised benefits and firearms context** while providing limited detail on automation provisions, immunity clauses, or scope expansion mechanisms
3. **Parliamentary debate engaged minimally** with Schedule 2's implications—one senator mentioned Robodebt connection across 159 pages of Hansard
4. **Constitutional architecture changed hours before passage** through State law incorporation not examined during inquiry
5. **IGIS warning about automated systems** issued 5 days before passage had limited apparent impact on enacted framework
6. **International transparency standards** for automated government decision-making not reflected in Schedule 2's design

Questions for the Commission:

- Do these observations suggest transparency standards for automated decision-making legislation should be strengthened?
- Should Explanatory Memoranda be required to disclose Royal Commission findings relevant to new administrative frameworks?
- Should plain-language summaries be mandatory for complex technical provisions affecting individual rights?
- Should parliamentary processes include minimum scrutiny periods for legislation introducing novel automated powers?
- What mechanisms might ensure that specialist warnings from oversight bodies like IGIS receive adequate parliamentary attention before passage?
- Should substantive amendments to constitutional architecture (such as State law incorporation) be subject to separate inquiry and consultation processes?

The broader pattern:

This analysis of Schedule 2's passage complements earlier Parts examining:

- Definitional ambiguity in antisemitism frameworks (Part I)
- Institutional design lessons from Hannah Arendt (Part VI)
- Historical patterns in Australian administrative policy (Part VII)
- Detailed statutory analysis of automation provisions (Part X)
- Private sector involvement questions (Part XI)

Together, these analyses suggest that **well-intentioned responses to genuine threats may create institutional vulnerabilities when:**

- legislative transparency is limited,
- parliamentary scrutiny is minimal due to time constraints,
- specialist warnings go unheeded or unincorporated,
- frameworks replicate features of previously problematic systems, and
- constitutional changes occur without adequate examination

The Commission's opportunity:

The Commission's findings on these matters could strengthen future legislative processes by establishing clearer expectations for:

- Transparency and disclosure when Parliament considers frameworks affecting individual rights
- Minimum consultation periods for complex technical legislation
- Integration of Royal Commission findings into policy development
- Mechanisms ensuring specialist oversight body recommendations receive serious parliamentary consideration
- Separate examination of substantive constitutional architecture changes

Such reforms could help ensure that responses to serious threats like antisemitism achieve protective goals through frameworks that preserve rather than undermine the democratic accountability and rule of law on which all communities' long-term safety ultimately depends.

PART X: SCHEDULE 2 STATUTORY FRAMEWORK—DETAILED ANALYSIS

Introduction: Examining the Automated Assessment Architecture

Schedule 2, Part 2, Division 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* inserted sections 92E–92G into the *Australian Security Intelligence Organisation Act 1979* (ASIO Act).

This Part provides detailed analysis of these statutory provisions, examining:

- The powers they create
- The safeguards they include (or omit)
- Their relationship to concerns raised by the Robodebt Royal Commission
- Their potential scope and application

The analysis draws extensively on the statutory text itself. Where the submission interprets statutory language, it does so by reference to established principles of statutory construction and recent Royal Commission findings on similar frameworks.

Note on approach: This Part quotes statutory provisions extensively. While this creates a technical and lengthy analysis, the submission considers detailed textual examination essential for evaluating whether Schedule 2 incorporates lessons from Robodebt or replicates structural features that proved problematic in that context.

I. Section 92E: Authorising Automated Assessment Processes

The Core Delegation Provision

Section 92E(1) states:

> The Director-General may, in writing, arrange for the use, under the Director-General's oversight, of computer programs to take specified assessment action that may, or must, be taken by the Organisation under this Act.

Analysis of key elements:

"arrange for the use": This language permits not only direct ASIO deployment but potentially third-party provision of software or systems. Whether this authorises private sector involvement in intelligence assessment depends on how "arrange for" is interpreted in practice.

"computer programs": The statute does not define what qualifies as a computer program for these purposes, nor does it specify technical requirements, validation processes, or certification standards.

"under the Director-General's oversight": The nature and extent of required oversight is not specified. The Robodebt Royal Commission found that oversight requirements without specific statutory definition often prove insufficient in practice, particularly where overseeing officials lack technical expertise to assess algorithmic operations.

"specified assessment action": Defined by ministerial instrument under section 92F (discussed below). This means the scope of what can be automated is not fixed by Parliament but determined by executive decision.

Legal Effect of Automated Decisions

Section 92E(5) provides:

> Specified assessment action taken by the operation of a computer program under an arrangement under subsection (1) is treated, for all purposes, as specified assessment action taken by the Organisation.

Significance:

This provision grants algorithmic outputs the same legal status as human decisions made by ASIO. Once a computer program takes "specified assessment action," that action is treated as if ASIO itself made the determination.

Implications for accountability:

When automated processes are legally equivalent to human decisions, questions arise about:

- Who bears responsibility when automated systems err?
- How can affected individuals challenge decisions made by algorithmic processes?
- What mechanisms exist for detecting systematic errors in automated assessments?

Retrospective Substitution Authority

Sections 92E(6)–(8) state:

> **(6)** The Organisation may take specified assessment action (the substituted action) in substitution for specified assessment action the Organisation is treated as having taken under subsection (5) if the Organisation is satisfied that the specified assessment action taken by the operation of the computer program is not the correct or preferable action.
>
> **(7)** The substituted action takes effect on:
> (a) if the Organisation specifies the day on which the substituted action takes effect—that specified day; or
> (b) otherwise—the day of the specified assessment action taken by the operation of the computer program.
>
> **(8)** The day specified under paragraph (7)(a) may be earlier than the day the substituted action is taken but not earlier than the day of the specified assessment action taken by the operation of the computer program.

Analysis:

These provisions authorise ASIO to:

1. Identify an assessment made by automated process as incorrect or non-preferable
2. Substitute a different assessment in its place
3. Make that substituted assessment take effect retroactively—from the date of the original automated assessment

Practical implications:

Consider a scenario where:

- In 2024, an automated system assesses Person X as low security risk
- Person X obtains employment requiring security clearance based on that assessment
- In 2026, ASIO determines the 2024 automated assessment was "not correct or preferable"
- ASIO substitutes a high-risk assessment with effect from 2024

The question arises: What is the legal status of decisions made in reliance on the original assessment? The statute does not address this temporal instability.

Comparison to Robodebt:

The Robodebt Royal Commission criticised retrospective application of unlawful debt calculations. While Schedule 2's retrospective provisions serve a different purpose (correction rather than initial determination), they create similar questions about temporal stability of administrative decisions.

II. Section 92F: Ministerial Definition of Scope

Non-Legislative Instrument Control

Section 92F provides:

> **(1)** The Minister may make an instrument for the purposes of subsection 92E(4) (meaning of specified assessment action).
>
> **(5)** An instrument made under subsection (1) is not a legislative instrument.

Significance:

The Minister, rather than Parliament, determines what "specified assessment action" means—i.e., what types of intelligence assessments can be automated.

The instrument is explicitly excluded from being a legislative instrument, which means:

- No parliamentary disallowance procedure applies
- No registration requirements under the Legislation Act
- No automatic sunset provisions
- Limited public visibility unless Minister chooses to publish

Implications for scope expansion:

While initial policy intent may focus on specific contexts (such as firearms licensing), the statutory language permits the Minister to expand automated assessment to other categories of intelligence action without returning to Parliament for legislative authorisation.

Democratic accountability questions:

This design raises questions about whether decisions to expand algorithmic assessment processes—which affect individual rights and liberties—should require parliamentary approval rather than executive discretion alone.

III. Section 92G: Safeguards and Immunity Provisions

The "Reasonable Steps" Requirement

Section 92G(1)–(2) state:

> **(1)** The Director-General must take all reasonable steps to ensure that specified assessment action taken by the operation of a computer program under an arrangement under subsection 92E(1) is specified assessment action that the Organisation could validly take under this Act.
>
> **(2)** Without limiting subsection (1), the Director-General must do the things (if any) prescribed by the regulations for the purposes of this subsection.

Analysis:

Subsection (1) imposes a duty on the Director-General to take "all reasonable steps" to ensure automated actions are valid.

Subsection (2) contemplates that regulations may prescribe specific requirements, but uses permissive language ("if any"), meaning regulations are not mandatory.

The Robodebt precedent on "reasonable steps":

The Robodebt Royal Commission found that "reasonable steps" requirements without specific statutory definition proved difficult to enforce. Commissioner Holmes observed that officials believed they were taking reasonable steps even while implementing an unlawful scheme.

The question arises whether section 92G(1)'s "reasonable steps" requirement will provide meaningful constraint or, like similar provisions in other contexts, operate as aspirational language without practical enforcement mechanisms.

The Statutory Note

The Note after subsection (2) states:

> Specified assessment action may still be invalid even if subsections (1) and (2) are complied with.

Significance:

This Note acknowledges that compliance with the "reasonable steps" requirement does not guarantee validity. Even if the Director-General takes all reasonable steps and complies with any regulations, the automated assessment "may still be invalid."

This creates an unusual statutory structure:

- Officials must take reasonable steps to ensure validity (subsection 1)
- Regulations may prescribe additional requirements (subsection 2)
- But compliance with both does not ensure validity (Note)
- Yet non-compliance does not affect validity (subsection 3, below)

The Immunity Provision

Section 92G(3) states:

> A failure to comply with subsection (1) or (2) does not affect the validity of the specified assessment action taken by the operation of a computer program under an arrangement under subsection 92E(1).

Plain language interpretation:

This provision means that even if the Director-General fails to take "all reasonable steps" to ensure validity, and even if regulations are not complied with, the automated assessment remains legally valid.

Relationship between safeguards and immunity:

The statutory structure creates what might be described as theoretical safeguards:

- **Subsection (1)** imposes a duty (take reasonable steps)
- **Subsection (3)** removes consequences for breach of that duty (non-compliance does not affect validity)

Comparison to administrative law principles:

Traditional administrative law holds that procedural failures affecting fairness can render decisions invalid. Section 92G(3) appears to displace this principle in the context of automated intelligence assessments.

The Robodebt parallel:

The Robodebt Royal Commission found that immunity provisions prevented effective review of unlawful automated decisions. Commissioner Holmes stated:

> "The Commonwealth entered into a scheme to obtain revenue through invalid assessments that were immune from legal challenge except through time-consuming and difficult court proceedings."

Section 92G(3) codifies a similar structure: assessments may be invalid (per the Note) but remain legally effective (per subsection 3) despite procedural failures.

IV. What the Statutory Framework Does Not Include

Absence of Specified Transparency Requirements

The automation framework in sections 92E–92G does not include statutory requirements for:

1. Algorithmic explainability:

- No requirement to document decision-making logic
- No requirement to disclose risk models or weighting factors
- No requirement to identify data sources used in assessments

2. Notification to affected persons:

- No requirement that individuals be informed they've been assessed
- No requirement to explain how automated assessment operated
- No requirement to disclose what information was considered

(Note: Section 36(1)(d) of the ASIO Act, which previously required notification of adverse security assessments, was deleted by Schedule 2—discussed in Part VII)

3. Review mechanisms:

- No statutory right to request human review of automated assessment
- No requirement for independent audit of algorithmic systems
- No merits review process for challenging automated determinations

4. Accountability structures:

- No requirement to publish algorithmic impact assessments
- No mandatory reporting on automated assessment outcomes
- No requirement for independent oversight of algorithmic deployment

Comparison to Robodebt Royal Commission Recommendation 17.1

The Royal Commission recommended:

> "The Commonwealth should not pursue or continue the development of automated decision-making systems where there are known or likely deficiencies in the data or the systems are otherwise unlikely to comply with relevant legal requirements."

More specifically, the Royal Commission emphasised:

- Transparent explanation of how automated systems operate
- Clear review pathways for affected individuals
- Independent scrutiny of algorithmic decision-making
- Accountability when systems produce harmful outcomes

Schedule 2's relationship to these recommendations:

While the Commission's recommendations were not legally binding, they represented considered expert analysis of what safeguards automated government decision-making requires.

The question arises whether Schedule 2's framework—which does not statutorily require the elements the Royal Commission identified as essential—reflects engagement with those findings or proceeds on a different view of what safeguards are necessary.

V. Comparative Analysis: International Approaches

European Union AI Act

The EU's regulatory framework for artificial intelligence in government includes:

- Mandatory transparency about algorithmic systems affecting individuals
- Requirements for human oversight of high-risk AI applications
- Explicit rights for affected persons to receive explanations
- Independent regulatory bodies with enforcement authority
- Public disclosure of impact assessments

United Kingdom Algorithmic Transparency Standard

The UK framework requires government departments using algorithms to:

- Publish descriptions of algorithmic systems
- Explain what data is used and how
- Provide information about testing and evaluation
- Establish accountability mechanisms
- Enable public scrutiny

Schedule 2 in Comparative Context

Schedule 2's approach differs from these frameworks by:

- Not requiring publication of algorithmic methodology
- Not mandating notification to affected persons
- Providing statutory immunity rather than review rights
- Relying on "reasonable steps" rather than specific transparency requirements
- Minimising rather than maximising public visibility

The question this raises:

Whether Australia's approach reflects considered judgment that our context requires different safeguards, or whether Schedule 2 was developed without full engagement with international best practices in algorithmic governance.

VI. Potential Applications Beyond Initial Context

The Statutory Language Permits Broad Application

While policy statements have framed Schedule 2 in terms of firearms licensing efficiency, the statutory language is not so limited:

Section 92E(1) applies to:

> specified assessment action that may, or must, be taken by the Organisation under this Act

"This Act" means the entire ASIO Act, not merely firearms-related provisions.

Section 92F(1) permits the Minister to define:

> specified assessment action

—without statutory limits on what categories of assessment may be included.

The Explanatory Memorandum acknowledges broad scope (paragraph 378):

> "Section 92G provides for oversight measures and safeguards on the use of automated decision-making for specified assessment action **(for example, assessments under Parts IV and Part IVA of the ASIO Act)**."

The phrase "for example" indicates these are illustrative, not exhaustive.

Mission Creep and Scope Expansion

Historical experience with administrative frameworks suggests that systems introduced for specific purposes often expand over time:

- Offshore detention facilities (established for Tampa crisis, operating 24 years later)
- Anti-terrorism laws (introduced post-9/11, applied to various contexts including journalism)
- Data retention frameworks (established for serious crime, accessed for minor matters)

The question for Schedule 2:

Whether safeguards exist to prevent automated intelligence assessment, introduced for specific contexts, from expanding to broader categories of security determination without returning to Parliament for explicit authorisation.

The ministerial instrument mechanism (section 92F) permits such expansion without legislative amendment—raising questions about whether this design adequately preserves parliamentary control over the scope of algorithmic government decision-making.

VII. Application to the Commission's Mandate

Relationship to Social Cohesion

The Commission's terms of reference include examining threats to social cohesion. Automated intelligence assessment frameworks affect social cohesion through several pathways:

1. Trust in government institutions:

When algorithmic systems operate with limited transparency, affected communities may lose confidence that decisions are made fairly or can be effectively challenged.

2. Perceived targeting of communities:

If automated assessment is perceived as disproportionately affecting particular ethnic, religious, or political communities, this may increase inter-community tensions rather than cohesion.

3. Chilling effects on legitimate activity:

Knowledge that automated surveillance or assessment occurs without notification may discourage lawful political participation, religious practice, or community organising.

4. Democratic legitimacy:

Systems perceived as having been introduced without adequate parliamentary scrutiny or public debate may face resistance from affected populations, undermining implementation effectiveness.

Relationship to Antisemitism Response

Schedule 2 was included in legislation responding to antisemitic violence. However, the provisions themselves are not limited to antisemitism-related assessments—they establish general-purpose infrastructure for automated intelligence processes.

The question this raises:

Whether creating broad algorithmic assessment infrastructure in response to specific threats serves protective goals, or whether more targeted, transparent, and accountable responses would better achieve both security and social cohesion objectives.

Conclusion: Questions for Systemic Reform

This detailed statutory analysis reveals that Schedule 2:

Creates substantial powers:

- Automated intelligence assessment substituting for human judgment
- Ministerial control over scope via non-legislative instrument
- Retrospective substitution of assessments with backdated effect
- Statutory immunity from invalidity despite safeguard failures

Lacks specified safeguards:

- No statutory transparency requirements
- No notification obligations to affected persons
- No explicit review mechanisms
- No independent oversight mandates

Differs from international approaches:

- EU and UK frameworks emphasise transparency and review rights
- Schedule 2 emphasises immunity and executive discretion

Was passed 28 months after Robodebt Royal Commission warned against:

- Automated decision-making without clear safeguards
- Immunity provisions preventing effective review
- "Reasonable steps" requirements as sole protection
- Systems affecting individuals without notification or appeal rights

The central question:

Does Schedule 2's architecture reflect considered engagement with Robodebt's lessons and international best practices, resulting in a framework appropriate for Australian context?

Or does it replicate structural features—automation, immunity, limited review, minimal transparency—that have proven problematic in other contexts?

This question is particularly relevant given that Schedule 2 was introduced in legislation responding to genuine security threats (antisemitic violence). **Well-intentioned responses to real dangers can nonetheless create institutional vulnerabilities if design choices do not adequately incorporate lessons from documented failures.**

The Commission's findings on these matters could inform whether Schedule 2 should be:

- Implemented as drafted
- Reformed to incorporate additional safeguards
- Replaced with alternative frameworks better balancing security and accountability

The detailed statutory analysis in this Part provides the foundation for assessing which path best serves both protective goals and democratic governance principles.

PART XI: PRIVATE SECTOR INVOLVEMENT IN AUTOMATED ASSESSMENT

Introduction: Examining Potential Third-Party Deployment

Part X analysed the statutory framework established by sections 92E–92G. This Part examines questions about potential private sector involvement in automated intelligence assessment—specifically, whether the phrase "arrange for the use" in section 92E(1) authorises contracting with third-party technology providers.

The analysis focuses on:

- The statutory language permitting such arrangements
- Recent developments in security clearances for technology companies
- International precedents for private sector involvement in intelligence automation
- Questions about transparency and accountability when proprietary systems are deployed

Note on evidence: This Addendum relies on publicly reported information about security clearances, government contracts, and company statements. Where information derives from media reporting or company disclosures, sources are identified.

I. The Statutory Language Regarding Third-Party Systems

"Arrange for the Use" — Textual Analysis

Section 92E(1) states:

> The Director-General may, in writing, arrange for the use, under the Director-General's oversight, of computer programs to take specified assessment action...

Interpretive questions:

The phrase "arrange for the use" could be interpreted as:

1. **Narrow interpretation**: ASIO develops software in-house, then "arranges" for its deployment
2. **Broad interpretation**: ASIO contracts with external vendors who provide algorithmic systems

Comparison to alternative statutory formulations:

If Parliament intended to require in-house development, alternative language might state:

- "develop and deploy computer programs"
- "use computer programs developed by the Organisation"

- "arrange for internal use of computer programs"

The use of "arrange for" without qualification appears consistent with the broader interpretation—permitting procurement from external providers.

Oversight Requirement

Section 92E(1)'s requirement that systems operate "under the Director-General's oversight" does not appear to restrict procurement source, only to require ASIO supervision regardless of origin.

Questions this raises:

- What constitutes adequate oversight of proprietary third-party systems?
- Can ASIO effectively oversee algorithms developed externally without access to source code?
- Does "oversight" require technical understanding sufficient to evaluate algorithmic operation?

II. Recent Developments: Security Clearances for Technology Providers

Palantir Technologies Security Clearance (November 2025)

Public reporting:

According to media reports (Michael West Media, 19 January 2026), Palantir Technologies received Protected-level security clearance under the Australian Information Security Registered Assessors Program (IRAP) in November 2025.

Significance:

Protected-level clearance authorises handling of information classified as PROTECTED, which includes:

- intelligence assessments,
- security evaluations, and
- law enforcement data.

Temporal relationship to Schedule 2:

- Palantir clearance: November 2025
- Antisemitism Bill introduced: 20 January 2026
- Interval: Two months

Interpretive questions:

Whether this timing reflects:

1. Coincidental administrative processes operating on separate timelines
2. Coordinated preparation for potential deployment under anticipated legislation

3. Normal procurement planning where clearances precede contract opportunities

The submission does not assert knowledge of internal government decision-making processes. However, the temporal proximity raises questions about whether security clearances were granted in anticipation of new statutory authorities.

What Protected-Level Clearance Permits

From Palantir's public statements (as reported):

The clearance enables Australian government agencies to use Palantir's:

- Foundry platform (data integration system)
- Artificial Intelligence Platform (AIP)

Capabilities described in company materials include:

- Real-time data integration from multiple sources
- Automated pattern recognition
- Predictive analytics and risk assessment
- Classification algorithms generating "confidence scores"

Relevance to sections 92E–92G:

These capabilities align closely with what "computer programs to take specified assessment action" (section 92E(1)) would require—suggesting technical compatibility between Palantir's systems and Schedule 2's framework.

III. Palantir's Operations in Other Jurisdictions

International Context

Purpose of this section:

Understanding how technology providers operate in other jurisdictions helps assess potential Australian deployment scenarios. This section summarises publicly reported information about Palantir's international operations.

United States: Immigration Enforcement Systems

Reporting by 404 Media and other outlets indicates:

Palantir developed systems for U.S. Immigration and Customs Enforcement (ICE) that:

- Generate detailed profiles of enforcement targets
- Integrate data from multiple databases
- Assign location probability scores
- Facilitate operational planning

Analytical relevance:

These systems demonstrate applied use of predictive algorithms for government decision-making affecting individuals—conceptually similar to automated security assessment, though in a different policy domain.

International Development Contexts

Public reporting on other deployments:

Palantir systems have been reported in use across multiple jurisdictions for various purposes including:

- Defence and intelligence operations
- Law enforcement analytics
- Border security applications

The analytical question:

Do international deployments indicate technical capacity and institutional experience relevant to Australian automated assessment contexts?

Corporate Statements and Public Positioning

CEO statements (publicly reported):

Palantir CEO Alex Karp has made public statements characterising the company's mission in direct terms, including references to defence and security applications.

Relevance to government procurement:

When governments procure technology from private vendors, corporate culture and public positioning may be relevant considerations for:

- Alignment with Australian values and legal frameworks
- Potential reputational implications
- Public confidence in outsourced systems

IV. Existing Australian Government Relationships

Contract History

Austender data indicates:

Palantir has received Australian government contracts totalling over $50 million since 2013, primarily across defence and national security agencies.

Reported 2024 revenue: $25.5 million from Australian operations

Significance:

These existing relationships indicate:

- Established commercial presence in Australian government sector

- Prior experience working with Australian agencies
- Institutional familiarity with Australian procurement processes

Lobbying Activities

Public reporting (Capital Brief, July 2025):

Palantir engaged CMAX Advisory (a government relations firm) following parliamentary scrutiny of the company's government contracts.

Timeline:

- July 2025: Engagement of lobbying firm
- November 2025: Protected-level security clearance granted
- January 2026: Antisemitism Bill introduced

Questions raised:

Whether lobbying activities influenced:

- Security clearance decisions
- Legislative design choices
- Procurement considerations

The submission does not assert improper influence. However, transparency principles suggest that relationships between lobbying, clearances, and legislative frameworks warrant public examination.

V. Statutory Framework and Private Sector Deployment

What Schedule 2 Permits Regarding Third-Party Systems

If ASIO "arranges for the use" (section 92E(1)) of Palantir or similar vendor systems:

1. Algorithmic assessments would have legal status

Section 92E(5): Automated action "is treated, for all purposes, as specified assessment action taken by the Organisation"

Implication: Third-party algorithmic output becomes official ASIO determination

2. Immunity would apply to vendor systems

Section 92G(3): Non-compliance with safeguards "does not affect the validity" of automated assessments

Implication: Proprietary algorithms immune from invalidation even if defective

3. Notification would not be required

Section 36(1)(d) deleted by Schedule 2

Implication: Individuals would not know third-party systems assessed them

4. Scope could expand via ministerial instrument

Section 92F: Minister defines "specified assessment action" via non-legislative instrument

Implication: Third-party deployment could expand without parliamentary approval

What Schedule 2 Does Not Require Regarding Third-Party Systems

The statutory framework does not include requirements for:

1. Transparency about vendor involvement:

- No requirement to disclose that third-party systems are used
- No requirement to identify which vendor provides algorithms
- No requirement to explain how vendor systems operate

2. Access to proprietary methodology:

- No requirement that source code be disclosed to ASIO
- No requirement that algorithmic logic be auditable
- No protection against vendor claims of trade secrecy

3. Data governance:

- No restriction on what data vendors can access
- No limitation on vendor data retention
- No prohibition on vendor use of Australian data for system training

4. Independent oversight:

- No requirement for independent audit of vendor systems
- No requirement that IGIS receive vendor methodology
- No enforcement mechanism against vendor non-compliance

Comparison to Public Sector Development

If ASIO developed systems in-house:

- Source code would be government property
- Methodology could be internally audited
- IGIS could theoretically access technical documentation
- Public accountability mechanisms might apply through normal government transparency processes

If ASIO procures vendor systems:

- Source code remains proprietary
- Methodology may be trade secret
- IGIS access depends on contract terms and vendor cooperation
- Trade secrecy claims may limit transparency

The question:

Whether procurement of proprietary algorithmic systems creates accountability gaps not present in public sector development scenarios.

VI. International Comparisons: Regulation of Private Sector AI in Government

European Union Approach

The EU AI Act addresses private sector involvement by:

- Requiring transparency about algorithmic systems regardless of provider
- Mandating conformity assessments for high-risk applications
- Establishing liability frameworks for both deployers and providers
- Requiring disclosure of system capabilities and limitations

United Kingdom Framework

UK guidelines for algorithmic procurement emphasise:

- Transparency about use of third-party systems
- Independent testing and evaluation
- Clear accountability structures
- Public disclosure of algorithmic impact assessments

Schedule 2 in Comparative Context

Australia's approach differs by:

- Not requiring disclosure of vendor involvement
- Not mandating independent evaluation of proprietary systems
- Providing statutory immunity rather than accountability mechanisms
- Limiting rather than enhancing transparency about algorithmic deployment

The analytical question:

Whether Australia's framework adequately addresses risks specific to private sector algorithmic deployment, or whether additional safeguards used in other jurisdictions might be relevant.

VII. Questions for Democratic Accountability

Transparency Concerns

If automated assessment uses proprietary third-party systems:

Affected individuals would not know:

- That algorithms assessed them (notification removed)

- Which company provided the system
- What data sources were consulted
- What risk factors triggered assessment
- What confidence scores were assigned

They would only learn:

- Adverse outcome occurred (clearance denied, application refused)
- Without explanation of automated process
- Without ability to challenge algorithmic determination (immunity provision)

Questions this raises:

- Does this level of opacity align with administrative law principles of procedural fairness?
- How can individuals exercise review rights if they don't know automated systems assessed them?
- What accountability exists when private sector errors affect government decisions?

Oversight Challenges

If IGIS seeks to oversee vendor systems:

Potential obstacles include:

- Trade secrecy claims limiting algorithm disclosure
- Lack of technical expertise to evaluate proprietary systems
- No statutory authority over private contractors
- No enforcement mechanisms against vendors

Questions this raises:

- Can IGIS effectively oversee systems it cannot fully examine?
- Should statutory frameworks require vendor cooperation with oversight bodies?
- What happens when vendor interests conflict with transparency?

Parliamentary Control

Under section 92F:

- Minister defines scope via non-legislative instrument
- Parliament cannot disallow ministerial determinations
- Vendor deployment could expand without legislative approval

Questions this raises:

- Should decisions to expand algorithmic assessment require parliamentary authorisation?
- Is executive discretion appropriate for frameworks involving private sector systems?
- How does Parliament maintain control over policy when implementation is outsourced?

VIII. Alternative Approaches: Safeguards for Third-Party Deployment

If Private Sector Involvement Proceeds

Should additional safeguards be considered:

1. Mandatory disclosure:

- Requirement to notify Parliament when third-party systems are deployed
- Public reporting of vendor identity and contract terms
- Disclosure of system capabilities and limitations

2. Technical transparency:

- Government ownership of source code or access rights
- Independent technical audit of vendor algorithms
- Requirement for explainable AI rather than black-box systems

3. Enhanced oversight:

- IGIS statutory authority to access vendor methodology
- Independent evaluation of proprietary systems before deployment
- Regular auditing of third-party algorithmic performance

4. Accountability mechanisms:

- Clear liability frameworks for vendor errors
- Review rights for individuals affected by automated assessments
- Enforcement procedures against non-compliant vendors

5. Scope limitations:

- Legislative specification of permitted applications
- Parliamentary approval required for scope expansion
- Sunset clauses requiring periodic reauthorisation

International Best Practice Models

Elements from EU and UK frameworks that might be adapted:

- Algorithmic impact assessments required before deployment
- Conformity testing for high-risk applications
- Public registers of algorithmic systems in use
- Independent regulatory oversight with technical expertise
- Clear liability and redress mechanisms

IX. Connection to Social Cohesion and Commission Mandate

Public Trust in Automated Systems

The Commission's social cohesion mandate connects to these questions through:

1. Perceived legitimacy:

When algorithmic systems operate with limited transparency—particularly if privately developed—affected communities may question whether assessments are fair, accurate, or subject to bias.

2. Accountability expectations:

If private sector errors affect government decisions, and immunity provisions prevent challenge, this may erode confidence that democratic institutions operate fairly.

3. Differential impact:

If automated systems disproportionately affect particular communities (whether intentionally or through algorithmic bias), this creates social division rather than cohesion.

4. Democratic participation:

Knowledge that opaque private sector systems assess citizens without notification may chill legitimate political participation, community organising, or religious practice.

Specific Concerns for Protected Communities

For Jewish Australians (the legislation's intended beneficiaries):

- Will private sector systems accurately identify genuine antisemitic threats?
- Or will algorithmic false positives direct resources away from real dangers?
- Does outsourcing assessment to vendors with international portfolios create security risks?

For communities potentially subject to assessment:

- Will algorithmic systems incorporate biases from training data or international deployments?
- Can affected individuals challenge errors when methodology is proprietary?
- Does lack of transparency increase perception of targeting, even if assessment is neutral?

Conclusion: Questions Requiring Consideration

This Addendum has examined whether Schedule 2's statutory language permits private sector involvement in automated intelligence assessment, and what questions such involvement raises.

Key observations:

1. **Statutory language appears to permit third-party procurement** ("arrange for the use")
2. **Recent security clearance for major technology provider** occurred two months before legislation introduced

3. **International precedents exist** for private sector algorithmic systems in government security applications
4. **Statutory framework lacks specific safeguards** addressing third-party deployment risks
5. **Accountability mechanisms are weaker** for proprietary systems than public sector development
6. **International frameworks include safeguards** not present in Schedule 2

Questions for the Commission:

- Should legislation explicitly address private sector involvement in automated assessment?
- Should additional transparency requirements apply to vendor systems?
- Should oversight bodies have statutory access to proprietary methodology?
- Should parliamentary approval be required before third-party deployment?
- What accountability mechanisms should exist when vendor errors affect government decisions?

The broader concern:

Automated decision-making affecting individual rights raises accountability questions regardless of whether systems are developed publicly or procured privately.

However, **private sector deployment introduces additional complexity:**

- Trade secrecy limiting transparency
- Proprietary interests potentially conflicting with accountability
- International vendor portfolios raising security and bias concerns
- Limited public sector technical capacity to evaluate external systems

The temporal context:

The proximity between Palantir's security clearance (November 2025) and Schedule 2's introduction (January 2026) raises questions about whether legislative design anticipated specific vendor deployment or was developed independently.

The Commission's consideration:

Whether Schedule 2 should be:

- Implemented as drafted (permitting vendor deployment without additional safeguards)
- Reformed to include transparency and accountability requirements for third-party systems
- Limited to public sector development until adequate safeguards for vendor involvement can be established

The answers to these questions affect whether automated intelligence assessment achieves security goals while maintaining democratic accountability and social cohesion—or whether outsourcing to private sector systems creates new vulnerabilities that undermine both objectives.

Part XII Additional AI-Specific Risks in Intelligence Assessment Contexts

Operational Consequences of Analytical Errors

Beyond the institutional dynamics already discussed, automated or AI-assisted analytical systems introduce specific technical risks relevant to intelligence assessment frameworks.

Unlike traditional bureaucratic processes where errors typically affect individuals sequentially, algorithmic systems can propagate incorrect assessments across multiple operational contexts simultaneously. Where security determinations flow into various administrative and operational systems, analytical errors may produce cascading practical consequences before detection occurs.

Potential operational scenarios include:

1. Transport and access control systems:

If automated security assessments feed into border control, airport screening, or facility access systems, incorrect risk classifications may result in:

- enhanced screening procedures applied based on erroneous intelligence
- access delays or denials affecting employment, travel, or services
- visible security responses that themselves create reputational harm

Such consequences may occur automatically once a classification enters operational systems, without human review of the underlying analytical basis.

2. Resource allocation effects:

Security and law enforcement resources are finite. Automated systems that generate false-positive risk assessments may:

- direct investigative attention toward individuals or groups not presenting genuine threats
- divert resources from actual security concerns
- create operational inefficiency through processing of unfounded alerts

3. Secondary administrative consequences:

Security assessments often influence decisions across multiple government functions:

- employment security clearances
- visa and citizenship determinations
- contract eligibility
- professional licensing

Where automated assessments propagate across these systems, a single analytical error may affect numerous administrative outcomes simultaneously.

The "Hallucination" Problem in AI-Assisted Analysis

Contemporary AI systems, particularly those using large language models or complex pattern-matching algorithms, can generate outputs that appear analytically coherent while being factually incorrect or unsupported by underlying data.

This phenomenon—often termed "hallucination" in technical literature—occurs when:

- systems extrapolate beyond available evidence
- pattern-matching algorithms identify correlations that are statistically spurious
- probabilistic models generate high-confidence outputs based on incomplete information

In intelligence contexts, this creates specific risks:

1. Plausibility without accuracy:

AI-generated assessments may read as professionally-written intelligence reports, complete with:

- confident analytical language
- apparent logical structure
- citations to data sources

Yet the substantive conclusions may rest on algorithmic artifacts rather than verified intelligence.

2. Detection challenges:

Unlike obvious system errors (crashes, null outputs, formatting failures), hallucinated content appears valid at surface level. Detection requires:

- technical expertise to audit algorithmic reasoning
- access to underlying data and model parameters
- time-intensive review of individual assessments

3. Scale effects:

Manual review becomes increasingly impractical as automated systems process larger volumes. The efficiency gains that make automation attractive also create challenges for quality assurance.

Professional Incentives in Classified Environments

An additional structural consideration concerns how security clearances and institutional hierarchy may affect error recognition.

Access to classified intelligence is conditional on clearance status, which itself depends on institutional trust. In such environments:

Personnel may face implicit professional risks in questioning analytical products:

- Challenging intelligence assessments may be perceived as questioning institutional competence
- Raising concerns about analytical tools may seem to imply inadequate vetting occurred
- For those responsible for deploying or maintaining systems, acknowledging errors may carry career implications

Senior officials may face reputational pressure as errors become visible:

- Admitting systematic flaws in deployed systems involves institutional exposure
- Political or public attention to intelligence failures may create pressure to minimise rather than transparently address issues
- The longer errors persist undetected, the greater the perceived cost of acknowledgment

This is not an allegation of misconduct. It describes ordinary human responses to professional and institutional risk in high-stakes environments.

The structural concern is:

If professional incentives discourage early surfacing of analytical uncertainty, and if secrecy limits external oversight, errors in automated intelligence assessment may persist longer than in open administrative systems.

By the time problems become undeniable, operational and human consequences may already be substantial.

Magnitude and Reversibility Concerns

The submission acknowledges that these concerns may appear speculative. However, recent Australian experience with Robodebt demonstrates that:

1. **Automated administrative systems can affect hundreds of thousands of people before errors are fully recognised**
2. **Reversing systemic errors once embedded in multiple downstream processes is extraordinarily difficult**
3. **Human costs—reputational harm, employment consequences, mental health impacts—may not be fully reversible even after errors are corrected**

If analogous dynamics were to occur in classified intelligence environments where:

- affected individuals have limited visibility into assessment processes
- external oversight faces structural constraints
- professional incentives may discourage early error acknowledgment

The potential for significant, difficult-to-reverse harm exists—not through malice, but through the structural characteristics of automated systems operating in high-authority, low-transparency environments.

The Efficiency Paradox: When Correction Costs Exceed Automation Benefits

A further practical consideration concerns the stated rationale for automated intelligence assessment: increasing processing capacity and reducing personnel burden.

This efficiency justification assumes that automated systems will operate with sufficient accuracy that correction requirements remain minimal. However, if structural factors discussed above—hallucinations, feedback loops, delayed detection in classified environments—result in significant error rates, the administrative burden of correction may exceed any efficiency gains.

Potential cost implications include:

1. Retrospective review requirements:

If systematic errors are eventually detected, agencies may need to:

- manually review all assessments produced during the affected period
- reconstruct evidentiary basis for decisions where records were overwritten
- contact affected individuals to correct classifications
- reverse administrative decisions based on erroneous assessments

The Robodebt experience demonstrated that retrospective correction can require years of work by hundreds of personnel—far exceeding the staffing "saved" through automation.

2. Litigation and review costs:

Where automated errors affect legal rights or entitlements:

- affected individuals may seek judicial review
- class actions may consolidate claims
- government legal resources must respond to challenges
- compensation frameworks may be required

These costs—both financial and reputational—can substantially exceed the administrative expenses automation was designed to reduce.

3. Institutional credibility damage:

If intelligence assessments lose perceived reliability due to known systematic errors:

- operational partners (law enforcement, allied agencies) may discount outputs
- courts may question evidentiary weight of intelligence products
- public confidence in security institutions may erode

Restoring institutional credibility after systematic failures can require decades, as ongoing challenges to historical offshore detention policy demonstrate.

4. Opportunity costs:

Personnel assigned to error correction, litigation response, and system remediation cannot simultaneously:

- conduct new intelligence analysis
- respond to emerging threats
- maintain operational capabilities

The resource diversion may leave agencies less capable than before automation was introduced.

The Robodebt Precedent on Correction Costs

The Robodebt scheme was justified as improving efficiency in welfare compliance. The subsequent costs included:

- **$1.8 billion in refunds and compensation** (far exceeding any savings)

- **Royal Commission costs** (extensive legal and administrative resources)
- **Ongoing legal proceedings** (individual cases still being resolved)
- **Institutional damage** (lasting impact on public trust in government administration)
- **Political consequences** (contributing to electoral outcomes and ministerial changes)

The administrative burden of attempting to correct systematic errors—identifying affected individuals, reconstructing accurate determinations, processing refunds, responding to legal challenges—consumed exponentially more resources than the automation "saved."

If similar dynamics occur with automated intelligence assessment:

- The efficiency justification disappears entirely
- Agencies face greater administrative burden than manual processes
- Human and institutional costs compound over time
- The supposed solution becomes a more expensive version of the original problem

The Irrecoverable Case

A particularly concerning scenario arises if:

1. Automated systems operate for extended periods before errors are detected (due to secrecy and professional incentive structures)
2. Original evidentiary records are progressively overwritten or obscured by derivative automated assessments
3. Error discovery occurs only after significant operational and human consequences have materialized

In such circumstances, **reconstruction of accurate assessments may become impossible**:

- Original data sources no longer accessible
- Multiple layers of automated analysis obscure provenance
- Affected individuals' circumstances have changed (employment lost, relationships damaged, opportunities foreclosed)
- Administrative decisions based on assessments have cascaded through multiple systems

At that point, agencies face a choice between:

- Maintaining known-to-be-unreliable assessments (because correction is impractical)
- Abandoning the entire dataset and starting fresh (rendering prior work worthless)

Neither option recovers the resources "saved" through automation. Both represent catastrophic failure of the efficiency justification.

Design Implication

The efficiency case for automated intelligence assessment depends entirely on accuracy rates being sufficiently high that correction requirements remain manageable.

Where structural factors—secrecy, professional incentives, AI hallucinations, feedback loops, delayed detection—may reduce accuracy or delay error discovery, the efficiency calculation reverses:

Automation becomes more expensive than manual processes because it generates errors at scale that require exponentially greater resources to correct than would have been needed to process assessments accurately in the first instance.

This is not a theoretical concern. **Robodebt demonstrated this exact dynamic** in an open administrative environment with relatively robust oversight.

In classified intelligence environments where:

- error detection is structurally delayed
- correction is more difficult
- affected individuals have limited recourse
- institutional incentives may favour minimization over transparency

The risk of efficiency-destroying error accumulation is substantially higher, not lower.

For this reason, the prudent approach—from a resource management perspective alone, setting aside human rights considerations—is to ensure robust safeguards are embedded before deployment, rather than assuming errors can be corrected after the fact.

Once automated systems operate at scale in classified environments, correction may be practically impossible regardless of resources committed.

Design Implication

The question raised is not whether AI-assisted tools should ever be used in intelligence contexts, but whether current safeguards are adequate to:

- detect and correct errors before they cascade through operational systems
- enable affected individuals to challenge assessments they cannot see
- ensure professional incentives favour transparency over institutional risk-management

Without such safeguards, even well-intentioned deployment of automated intelligence assessment tools may create vulnerabilities that exceed those present in traditional human-directed analytical processes.

The Oversight Challenge: When Performance Reporting Is Also Automated

A further structural consideration arises where oversight mechanisms themselves may be affected by automation.

Traditional oversight of intelligence agencies—including parliamentary committees, IGIS, and Royal Commissions—relies substantially on agencies' good-faith reporting about operational performance. ASIO and other agencies have demonstrated responsiveness to oversight bodies and have provided detailed information about how expanded powers granted post-2001 have been exercised.

However, the introduction of AI-assisted systems, particularly those provided by external vendors, may create new oversight challenges that existing mechanisms are not designed to address.

Automated performance reporting:

AI systems can generate performance reports automatically, including:

- statistics on assessments processed
- accuracy metrics based on internal validation
- efficiency comparisons to manual processes
- compliance documentation

Such reports can be produced in minutes and may appear comprehensive and professionally formatted.

The structural challenge:

If the same automated system that makes assessments also generates reports on its own performance:

- Self-assessment may not identify systematic biases or errors the system is designed to overlook
- Performance metrics may reflect what the system is optimized to measure, not necessarily operational accuracy
- Generated reports may appear authoritative while missing categories of error the system cannot detect

Vendor reporting incentives:

Where third-party vendors provide AI systems under commercial contracts:

- Vendor financial interests favor demonstrating system success
- Commercial incentives may not align with transparent disclosure of limitations
- Proprietary systems may limit independent verification of vendor claims
- Contract renewal and expansion depend on positive performance assessment

This does not allege vendor dishonesty. It describes ordinary commercial dynamics where parties report information that supports their interests.

The oversight gap:

Traditional oversight mechanisms assume:

- human officials reporting on human decision-making processes
- institutional incentives for accuracy (career reputational stakes)
- ability to interview decision-makers and review reasoning processes

With automated systems and external vendors:

- "Decision-makers" are algorithms, not interviewable humans
- Vendor reporting may be optimized for commercial rather than accountability purposes
- Oversight bodies may lack technical capacity to independently audit AI system performance
- Performance reports can be generated faster than they can be meaningfully reviewed

Implications for Royal Commission oversight:

If ASIO reported to this Commission that automated systems were performing well:

- Would the Commission have independent capacity to verify claims?
- Would vendor-provided performance data be auditable by non-technical reviewers?
- Could oversight identify systematic errors if the reporting system itself is automated?
- Would commercial confidentiality limit access to system internals?

Recent precedent:

The Robodebt experience included officials reporting to Parliament that the system was operating lawfully and effectively—while internal legal advice warned otherwise. The gap between external reporting and internal reality persisted for years.

With AI systems and vendor involvement, similar gaps could emerge more easily because:

- Performance reporting can be automated (reducing human awareness of problems)
- Vendor interests favor positive reporting (reducing transparency about limitations)
- Technical complexity creates barriers to independent verification (reducing oversight capacity)

Design Implication for Oversight

For these reasons, safeguards for AI-assisted intelligence assessment should include:

1. Independent technical audit capacity:

- Oversight bodies require access to qualified technical experts
- Authority to audit algorithmic systems, not just receive vendor reports
- Capacity to test system performance against ground truth

2. Separation of performance reporting from system operation:

- Independent validation of accuracy claims
- Human-conducted review of automated performance metrics
- External verification of vendor-provided assessments

3. Transparent disclosure of limitations:

- Statutory requirement that system limitations be reported alongside performance metrics
- Disclosure of error rates, false positives, and correction requirements
- Documentation of instances where automated assessments were overridden

4. Enhanced IGIS authority regarding vendor systems:

- Statutory access to proprietary algorithmic methodology
- Authority to require independent technical evaluation
- Capacity to review vendor contracts and performance incentive structures

Without such mechanisms, oversight may become performative:

Agencies report what automated systems tell them. Vendors provide optimistic performance data. Oversight bodies lack capacity to verify claims independently. Problems remain undetected until operational failures become undeniable.

This is not a prediction that ASIO will act in bad faith or that vendors will deliberately mislead. It is an observation that existing oversight frameworks were designed for human decision-making processes and may not adequately address challenges posed by automated systems and commercial vendor involvement.

The Commission's consideration of whether current oversight mechanisms require strengthening before AI-assisted intelligence assessment is deployed could help prevent future accountability gaps similar to those documented in the Robodebt context.

The 2026 Census: A High-Risk Data Matching Event

The timing of sections 92E–92G's enactment creates a particular concern regarding the 2026 Australian Census, scheduled for August 2026.

The Census represents the most comprehensive population data collection exercise conducted by the Australian government, gathering information from every household including:

- household composition and relationships
- country of birth and ancestry
- languages spoken
- religious affiliation
- employment and education status
- address and location data

The Ethnicity Question and Profiling Risk:

The Census includes questions about ancestry and ethnic origin. While Census data is ostensibly anonymized, the combination of address-based collection and comprehensive demographic detail creates significant data matching potential when combined with other government datasets.

Data Matching Capabilities:

Modern algorithmic systems can readily match Census records (organized by address) with other government databases containing:

- driver license registries (linked to addresses and names)
- electoral rolls (names, addresses, demographic indicators)
- visa and citizenship records (nationality, country of origin)
- property ownership records (names, addresses)
- state and territory databases (various purposes, address-linked)

An AI system with access to multiple datasets could reconstruct individual identities from "anonymized" Census data by:

1. Using address as the linking key across datasets
2. Matching household composition patterns
3. Cross-referencing demographic indicators
4. Building individual profiles including ethnicity, religion, and ancestry

Once matched, the system could:

- Create population-wide profiles organized by ethnicity or religious background
- Identify communities or neighborhoods by demographic composition
- Flag individuals or households meeting particular profile criteria
- Feed these assessments into security evaluation frameworks

The Transparency Problem:

Unlike traditional intelligence gathering which requires human analysts to manually compile information, algorithmic data matching can occur:

- automatically and at scale
- without human review of individual cases
- without notification to affected individuals
- without public visibility of matching methodologies
- within classified systems exempt from Freedom of Information

The public would have no way of knowing whether:

- Census data had been matched with other datasets
- Their household's ethnicity data had been incorporated into security assessments
- Algorithmic profiling based on ancestry or religion had occurred
- Such profiling influenced administrative decisions affecting them

Connection to Sections 92E–92G:

The ASIO Act amendments create precisely the infrastructure that could enable Census data matching:

- **Section 92E(1)** authorizes use of computer programs for "specified assessment action"
- **Section 92F** allows ministerial instruments to define scope without parliamentary amendment
- **Section 92G(3)** provides immunity from invalidity even if processes fail
- **No notification requirement** means individuals would not know their data was matched
- **Classification** prevents public scrutiny of matching methodologies

The combination of:

1. Comprehensive Census data (including ethnicity)
2. Address-based collection enabling cross-dataset matching
3. Algorithmic assessment powers under sections 92E–92G
4. Lack of notification or transparency requirements
5. Immunity provisions limiting legal challenge

Creates the precise conditions for population-wide profiling based on ancestry, religion, or ethnic origin—without public knowledge or individual recourse.

This is not speculative:

- The technical capacity for such matching exists
- The legal authority (if sections 92E–92G remain) would exist
- The data sources (Census plus other databases) would exist

- The institutional infrastructure (ASIO algorithmic assessment) would exist
- The only missing element is the decision to implement it

Historical Precedent:

Census data has been misused for targeting minority populations in other democratic societies during periods of perceived security threat. The United States' use of Census data to identify and locate Japanese Americans for internment during World War II demonstrates that even democratic governments with constitutional protections can use population data for discriminatory purposes when institutional safeguards are absent.

The Timing Concern:

The August 2026 Census occurs shortly after sections 92E–92G's enactment and potentially during the Royal Commission's deliberative period. If these provisions remain operational and Census data matching occurs, reversal would be extraordinarily difficult:

- Data once matched cannot be "unmatched"
- Profiles created cannot be reliably deleted from all systems
- Downstream uses (assessments, decisions) cannot be fully traced
- Affected individuals cannot be comprehensively identified for notification

The Precautionary Approach:

Given these risks, this submission respectfully urges the Commission to recommend that:

1. **Census data be explicitly excluded** from automated intelligence assessment frameworks pending comprehensive review
2. **Specific statutory protections** be enacted prohibiting matching of Census data with security databases without:
    - Individual warrants based on specific evidence
    - Judicial authorization
    - Notification requirements (except where immediate operational necessity prevents it)
    - Regular public reporting (in aggregated form) of any matching conducted
3. **The 2026 Census be conducted** under enhanced privacy protections given the timing relative to new security assessment powers
4. **Independent audit mechanisms** be established to verify Census data is not being matched with other datasets for profiling purposes

The Alternative:

Without such safeguards, the 2026 Census could become the foundational dataset for algorithmic profiling of Australian residents based on ancestry and religion—creating precisely the classification systems this submission has argued pose fundamental risks to liberal democratic governance.

The Commission has an opportunity to recommend protective measures before the Census occurs, rather than investigating misuse after population-wide data has been collected and potentially matched.

Connection to Recommendation 17.1

The Robodebt Royal Commission's emphasis on transparency, reviewability, and independent oversight reflects precisely these concerns:

Where automated systems can:

- affect large populations rapidly
- create cascading consequences across multiple administrative domains
- operate with limited visibility to affected persons
- and face structural obstacles to early error correction

Then safeguards must include:

- Transparent explanation of how systems operate
- Clear review pathways for affected individuals
- Independent audit capacity to detect systematic errors
- Institutional incentives favouring early problem identification

These principles apply with particular force where automation enters classified environments, because the same factors that make intelligence work necessary—secrecy, authority, and limited external contestability—also amplify risks when automated analytical errors occur.

PART XIII: RECOMMENDATIONS TO THE ROYAL COMMISSION

Overview: A Tiered Framework

This submission respectfully proposes a **tiered framework of recommendations**, reflecting differing levels of proximity to the Commission's mandate and differing degrees of immediacy.

- **Tier 1 recommendations** address matters directly connected to Schedule 2 amendments to the ASIO Act, automated decision-making safeguards, and democratic accountability within the Commission's terms of reference.
- **Tier 2 recommendations** propose structural governance reforms that may help prevent similar accountability challenges in future legislative processes.

Introduction: From Analysis to Action

Parts I through XII have documented:

- questions about transparency in how Schedule 2 was presented to Parliament;
- concerns about whether the legislative process provided adequate opportunity for informed consideration;
- structural similarities between Schedule 2 and features the Robodebt Royal Commission identified as problematic;
- questions about whether lessons from Robodebt were adequately incorporated into Schedule 2's design; and
- the opportunity now before the Commission to recommend safeguards before widespread implementation occurs.

This Part translates that analysis into **specific, actionable recommendations**.

The Commission has authority to:

- make findings about what occurred;
- recommend legislative amendment;
- refer matters for investigation or prosecution;
- recommend institutional reforms; and
- set standards for future automated decision-making.

These recommendations are structured by **priority and feasibility**:

- **Tier 1 (Most important / urgent / achievable):** prevention and immediate accountability.
- **Tier 2 (Secondary / structural reforms):** reforms to ensure similar concerns do not recur.

RECOMMENDATION 1 (TIER 1): PRIMARY — REPEAL OR SUBSTANTIALLY AMEND SCHEDULE 2 CHANGES TO THE ASIO ACT (SECTIONS 92E–92G)

A. The Finding This Submission Urges

This submission respectfully urges the Commission to consider finding:

1. Sections 92E–92G of the Australian Security Intelligence Organisation Act 1979 (ASIO Act), as inserted by **Schedule 2** to the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*, were enacted following a legislative process that raises questions about adequate parliamentary scrutiny and informed consideration.
2. The Explanatory Memorandum may not have adequately disclosed material features including:
    - the relevance of the Robodebt Royal Commission and Recommendation 17.1;
    - the practical consequences of immunity from invalidity provisions;
    - the scope implications of "arrange for" authority;
    - retrospective substitution provisions;
    - the deletion of notification requirements; and
    - the full breadth of the algorithmic assessment framework's potential application.
3. Parliamentary debate suggests limited substantive engagement with the implications of Sections 92E–92G (including minimal discussion of algorithmic assessment scope, immunity effects, notification removal, and review pathways).
4. The Schedule 2 framework raises questions about consistency with each element of Robodebt Royal Commission Recommendation 17.1 (review pathways, transparency about algorithms, and availability for independent scrutiny).
5. The framework exhibits structural features that the Robodebt Royal Commission identified as problematic (automation affecting individuals, limited transparency, constrained review mechanisms, and diffusion of responsibility), enacted within 28 months of that Commission's report.

B. The Recommendation This Submission Urges

This submission respectfully urges the Commission to recommend:

Parliament should repeal Sections 92E–92G of the ASIO Act in their entirety and, if any automated intelligence assessment is to be used, replace them with a framework that:

1. prohibits immunity from invalidity;
2. requires notification before adverse action affecting individual rights;
3. mandates human decision-making (algorithms may advise, humans must decide);
4. ensures transparency of methodology (distinct from protection of classified sources);
5. provides independent merits review;

6. limits scope through primary legislation (Parliament defines categories, not ministerial instruments);
7. prohibits retrospective operation; and
8. addresses questions about private sector and foreign contractor involvement in core intelligence assessment functions.

A detailed illustrative alternative model is provided in **Part XIV** (with technical detail capable of being included as an Appendix).

C. The Political Sustainability Argument

This submission respectfully notes that the primary recommendation for repeal or substantial amendment rests not only on Robodebt parallels and procedural concerns, but on a fundamental institutional reality:

Legislative powers persist across changes of government.

The Commission may encounter arguments that sections 92E–92G can be safely implemented because current political circumstances favour appropriate use. This submission respectfully argues that such arguments fundamentally misunderstand the nature of institutional tools.

Statutory powers do not expire with political influence. They persist.

Australia's political landscape includes concerning far-right elements. This is documented reality, not speculation: individuals displaying Nazi symbols and antisemitic messaging have demonstrated at state parliament buildings. Political parties with established records of targeting minority communities based on religion, ethnicity, and immigration status are polling at levels that could translate to increased parliamentary representation and potential coalition influence.

Should such elements gain governmental influence—through electoral success, coalition arrangements, or influence on major party policy—the algorithmic assessment capabilities, ministerial discretion, and immunity provisions created by sections 92E–92G will be available for repurposing.

This is not speculation about unlikely scenarios. This is analysis of how institutional tools operate across political transitions.

The question is not whether current architects intended to protect the Jewish community. The question is what happens to these tools when very different people control them—whether in five years, ten years, or thirty.

The Commission may wish to consider whether recommending that Parliament maintain tools that future governments could weaponize against the very community the current government sought to protect represents sound institutional design.

This submission respectfully argues it does not.

Effective protection of the Jewish community requires democratic accountability and institutional safeguards, not reliance on the durability of current political alignments.

D. Interim Measure (Non-operationalisation)

Pending repeal, amendment, and constitutional advice, this submission respectfully urges the Commission to recommend that the Attorney-General advise the Director-General of ASIO and the Minister for Home Affairs that:

1. Sections 92E–92G should not be operationalised pending:
    - Solicitor-General advice on constitutional questions;
    - parliamentary reconsideration with full information disclosure; and
    - implementation of safeguards consistent with Recommendation 17.1.
2. No systems should be acquired or deployed and no contracts entered into under s 92E "arrange for" authority during this period.

RECOMMENDATION 2 (TIER 1): REFERRALS FOR INVESTIGATION AND REVIEW

This submission respectfully urges the Commission to consider recommending referrals or equivalent investigative steps as appropriate to the evidence.

A. NACC

Investigation into whether Explanatory Memorandum preparation involved questions about disclosure of material information; and whether the timing and sequence of contractor clearances and legislative development raise questions requiring examination.

B. IGIS

Investigation into ASIO's role in the development of, advice regarding, and institutional posture toward Sections 92E–92G and any associated contractor engagement pathways.

C. CDPP / AFP

If supported by evidence, assessment of whether conduct relating to the Explanatory Memorandum or parliamentary process warrants consideration under relevant offences (including false or misleading information, misconduct, or related offences).

D. APSC

Examination of potential APS Code considerations, questions about institutional learning from Robodebt, and whether accountability mechanisms require strengthening to prevent recurrence.

E. Solicitor-General Advice

Obtain advice addressing:

- Chapter III and section 75(v) implications of immunity provisions;
- procedural fairness considerations (including notification requirement removal);

- retrospectivity questions;
- scope delegation and ministerial expansion considerations; and
- any other constitutional vulnerabilities identified on review.

RECOMMENDATION 3 (TIER 2): MINIMUM LEGISLATIVE SAFEGUARDS IF ANY AUTOMATION FRAMEWORK REMAINS

If Parliament retains any form of automated or algorithmic decision-support in this domain, the Commission should recommend minimum safeguards:

1. **No immunity from invalidity** — administrative decisions must be subject to judicial review for lawfulness;
2. **Notification before adverse action** — individuals must be informed when algorithmic assessments affect their rights;
3. **Human decision-making preserved** — algorithms may assist analysis but humans must make final determinations;
4. **Transparency of methodology** — how systems operate must be public (distinct from protecting classified sources);
5. **Independent merits review** — structured appropriately for classified material contexts;
6. **Scope defined by primary legislation** — Parliament determines categories, not ministerial instruments;
7. **No retrospective operation** — assessments apply prospectively only; and
8. **Sovereignty controls** on private sector and foreign contractor involvement in core assessment functions.

(Detailed illustrative drafting options are provided in Part XIV to demonstrate technical feasibility.)

RECOMMENDATION 4 (TIER 2): PARLIAMENTARY PROCESS AND EXPLANATORY MEMORANDUM REFORMS

This submission respectfully urges reforms to strengthen transparency and informed consideration in future legislative processes, including:

- **Mandatory disclosure** where legislation touches Royal Commission subject matter (requiring explanation of how recommendations were considered);
- **Plain language disclosure** of immunity clauses and automated decision-making provisions in Explanatory Memoranda;
- **Minimum scrutiny periods** for complex legislation, particularly involving security matters and automated systems;
- **Strengthened PJCIS resourcing** and access to independent technical expertise; and

- **Sunset clauses with independent evaluation** for high-risk automated decision systems affecting individual rights.

RECOMMENDATION 5 (TIER 2): ROBODEBT RECOMMENDATION 17.1 COMPLIANCE MONITORING

Recommend:

- **Annual reporting to Parliament** on implementation of Robodebt Royal Commission recommendations across Commonwealth agencies;
- **Cross-Commonwealth audit** of automated systems for Recommendation 17.1 compliance; and
- **Transparency about Royal Commission referrals** — a structured mechanism for reporting on the status of sealed referrals (in appropriately aggregated form where investigations are active), to ensure Royal Commission findings receive appropriate follow-up.

RECOMMENDATION 6 (TIER 2): ESTABLISH ALGORITHMIC ACCOUNTABILITY OVERSIGHT

Recommend creation or expansion of oversight capacity to:

- **Register all Commonwealth automated decision-making systems** affecting individual rights;
- **Audit for fairness, accuracy, explainability, and bias** in algorithmic processes;
- **Require algorithmic impact assessments** for high-risk systems before deployment;
- **Receive complaints and require remediation** when systems produce problematic outcomes; and
- **Report annually to Parliament** on automated decision-making across government.

For intelligence agencies, recommend coordinated oversight between any such body and IGIS (IGIS retains primacy on intelligence operations; technical algorithmic accountability expertise complements operational oversight).

RECOMMENDATION 7 (TIER 1): PROTECT 2026 CENSUS DATA FROM ALGORITHMIC PROFILING

Given the timing of the August 2026 Census and the recent enactment of sections 92E–92G, this submission respectfully urges urgent recommendations to prevent Census data being used for population-wide profiling:

Immediate Safeguards Required:

1. **Statutory prohibition on Census data matching** with security databases for automated assessment purposes
2. **Explicit exclusion of Census data** from "specified assessment action" that may be defined under section 92F ministerial instruments
3. **Enhanced Census privacy protections** including:
    - Clear statutory prohibition on matching Census address data with driver license, electoral roll, or other identification databases for security profiling
    - Requirement that any use of Census data for individual security assessment requires judicial warrant based on specific evidence
    - Mandatory notification to affected individuals (except where immediate operational necessity prevents it)
    - Regular public reporting of any Census data access by intelligence agencies
4. **Independent audit mechanism** to verify Census data is not being matched with other datasets for profiling based on ancestry, religion, or ethnic origin

Rationale:

The 2026 Census will collect comprehensive data including ethnicity and religious affiliation. AI systems can readily match address-based Census records with other government databases (driver licenses, electoral rolls, property records) to reconstruct individual identities and create population-wide profiles.

The combination of:

- Comprehensive Census data (including ethnicity questions)
- Address-based collection enabling cross-dataset matching
- Algorithmic assessment powers under sections 92E–92G
- Lack of notification requirements
- Immunity provisions

Creates conditions for population-wide profiling based on ancestry or religion without public knowledge or individual recourse.

The timing is critical: The Census occurs in August 2026, shortly after sections 92E–92G's enactment. If data matching occurs, reversal becomes extraordinarily difficult. The Commission has an opportunity to recommend protective measures before the Census occurs.

Connection to Jewish Community Protection:

While Census data could theoretically be used to profile any community, the Jewish community has particular reason for concern. Historical precedents (including use of Census data to identify Japanese Americans for internment) demonstrate that even democratic governments can misuse population data during periods of perceived security threat.

Creating statutory protections for Census data protects all communities—including the Jewish community—from being identified and profiled based on ancestry or religious responses to Census questions.

RECOMMENDATION 8 (TIER 2): COMMUNITY CONSULTATION AND PUBLIC EDUCATION

Recommend that any automated decision-support frameworks in security contexts require:

- **Genuine consultation** with affected communities and civil liberties organisations prior to deployment; and
- **Independent public education** on what automated decision-making means, how it operates, and how review pathways function.

RECOMMENDATION 9 (TIER 2): INTERNATIONAL BENCHMARKING

Recommend an independent comparative review of international models (European Union, United Kingdom, Canada, New Zealand) and incorporation of best practice standards into Commonwealth governance of automated decision-support systems.

RECOMMENDATION 10 (TIER 2): IMPLEMENTATION PATHWAY

Recommend a phased implementation pathway:

- **Immediate (0–30 days):** Non-operationalisation advice; referrals initiated; Solicitor-General advice requested.
- **Short-term (30–90 days):** Parliamentary reconsideration or inquiry with adequate submissions period; IGIS investigation commenced.
- **Medium-term (3–12 months):** Legislative amendments enacted; oversight mechanisms established; Explanatory Memorandum and scrutiny reforms implemented.
- **Long-term (1–5 years):** Sustained audit regime; international best practice alignment; embedded safeguards monitored and refined.

Conclusion

These recommendations are structured to provide:

- **Immediate harm prevention and accountability** (Tier 1) — addressing Schedule 2's specific vulnerabilities before widespread implementation;
- **Durable institutional safeguards against recurrence** (Tier 2) — ensuring that lessons from Robodebt and Schedule 2 strengthen future legislative processes and automated decision-making frameworks.

Part XIV then provides an illustrative framework showing how algorithmic decision-support could operate lawfully and accountably — demonstrating that reconsideration of Schedule 2 does not require rejecting technology, but requires technology to operate under law with appropriate safeguards.

PART XIV — Alternative Framework for Accountable Algorithmic Governance

(Supporting Material — Non-Prescriptive Model)

This Appendix provides an **illustrative governance framework** demonstrating how algorithmic decision-support could operate within intelligence and national security contexts while remaining consistent with democratic accountability, constitutional principles, and Recommendation 17.1 of the Royal Commission into the Robodebt Scheme.

The material in this Appendix is **not presented as a final legislative proposal**, nor does it purport to prescribe a specific statutory model. Rather, it is included to assist the Commission by showing that alternatives to the current Schedule 2 changes to the ASIO Act are technically and institutionally feasible. The examples below are conceptual and explanatory in nature, intended to support analysis in the main submission by illustrating how safeguards such as human decision-making, transparency of methodology, independent review, and oversight architecture could operate in practice.

Nothing in this Appendix should be read as limiting the Commission's discretion or recommending any particular drafting approach. Its purpose is to demonstrate that rejecting or reconsidering Schedule 2 does not require abandoning technological tools; instead, it shows how such tools could be implemented in a manner consistent with rule-of-law principles and long-standing administrative safeguards.

Introduction: Demonstrating the Possible

This submission has documented extensive failures:

- Schedule 2's concealment and immunity provisions
- Robodebt's unlawful architecture replicated
- A 27-year pattern from Tampa through offshore detention to algorithmic automation
- Democratic accountability systematically foreclosed

This Part demonstrates that another path exists.

Algorithmic intelligence assessment is not inherently problematic. Technology can legitimately assist human decision-makers in processing large volumes of information, identifying patterns, and allocating investigative resources efficiently. Many democracies successfully use algorithmic tools in law enforcement and intelligence while maintaining robust accountability.

The question is not whether to use technology, but how to use it lawfully.

This Part respectfully offers the Commission a detailed framework showing what legitimate algorithmic intelligence assessment would look like - complying with Robodebt Royal Commission Recommendation 17.1, respecting constitutional limits, maintaining democratic accountability, and demonstrating that security and transparency are compatible, not contradictory.

This is not a theoretical exercise. This framework draws from:

- International best practices (UK, Canada, New Zealand, European Union)
- Australian precedents (existing automated systems with proper safeguards)
- Expert recommendations (Law Reform Commission, Ombudsman, IGIS)
- Technical feasibility (what current technology actually enables)

This submission suggests the Commission can show Parliament: refusing Schedule 2 does not mean refusing technology. It means requiring technology to serve justice, not evade accountability.

I. CORE PRINCIPLES FOR LEGITIMATE ALGORITHMIC GOVERNANCE

This submission respectfully suggests that any algorithmic decision-support system in the intelligence or law enforcement context must adhere to eight core principles:

Principle 1: Human Decision-Making Preserved

Algorithms may advise. Only humans may decide.

- Computer programs can process data, identify patterns, calculate risk scores
- **But**: Humans must make final determinations affecting individual rights
- No adverse security assessment without human review of algorithmic output
- No automated rejection of visa/citizenship without human consideration
- Human decision-maker must be identified, accountable, and trained

Rationale: Algorithms lack judgment, context, and accountability. Humans understand circumstances computers cannot process. When errors occur, humans can be held responsible; code cannot.

Robodebt precedent: Commissioner Holmes found automated decision-making without meaningful human review was unlawful and unjust.

International standard: European Union AI Act (2024) prohibits automated decision-making in high-risk contexts including law enforcement and immigration without human oversight.

Principle 2: Notification Before Adverse Action

Affected persons must know before being harmed.

- If algorithmic assessment contributes to visa denial, person must be notified before denial
- If assessment affects employment/travel/licensing, person notified before harm
- Notification must include: that algorithmic assessment was used, opportunity to respond, review mechanism
- Security sources can be protected while still providing notification

Rationale: Natural justice requires notice and opportunity to be heard. Secret adverse determinations are incompatible with rule of law.

Robodebt precedent: Failure to properly notify debt basis was central to unlawfulness finding.

Schedule 2 failure: Section 36(1)(d) deletion removes notification requirement entirely.

Principle 3: Methodology Transparency, Source Protection

How assessments work must be public. Who provided intelligence can remain classified.

This submission suggests distinguishing:

MUST be public (methodology):

- What data sources are considered (open source, law enforcement databases, intelligence reports)
- What factors are weighted (travel patterns, associations, communications, financial transactions)
- How risk scores are calculated (algorithms, weighting, thresholds)
- What triggers adverse assessment (score levels, specific combinations)
- How often reviewed and updated (annual, after incidents, independent audit)

CAN remain classified (sources):

- Specific intelligence reports from foreign partners
- Identities of human sources
- Technical collection capabilities
- Ongoing investigations

Example: "We assess risk based on travel to conflict zones, communications with listed organizations, and financial transactions with designated entities, weighted according to recency and frequency" = **transparent methodology**.

"The specific intelligence report from Five Eyes partner indicating this person traveled to Syria" = **classified source**.

Rationale: Methodology transparency enables meaningful review without compromising sources. Courts, AAT, IGIS, affected persons can assess whether methodology is lawful and reasonable without seeing classified intelligence.

Technical feasibility: UK MI5 publishes methodology for threat assessments while protecting sources. Canadian CSIS provides algorithmic methodology to Federal Court in camera for review. This is standard practice internationally.

Principle 4: Independent Review on the Merits

Courts and tribunals must be able to assess whether determinations are correct, not merely whether procedures followed.

- Administrative Appeals Tribunal (AAT) has jurisdiction to review adverse security assessments
- AAT can review in camera with security-cleared members

- AAT assesses whether determination was correct on evidence, not merely procedurally proper
- AAT can access classified material, receive confidential briefings, question decision-makers
- AAT decisions subject to judicial review on questions of law

Rationale: Procedural review alone (did ASIO follow its internal processes?) is insufficient. Merits review asks: was the determination actually justified by evidence?

Precedent: AAT reviews many security-sensitive matters (visa cancellations, citizenship denials, access to classified information) with appropriate protections. This is routine, not radical.

Schedule 2 failure: Section 92G(3) immunity clause prevents meaningful review by removing courts' ability to invalidate unlawful assessments.

Principle 5: No Immunity from Invalidity

Unlawful determinations must be voidable.

- If algorithmic assessment violates law, courts can declare it invalid
- If methodology is discriminatory, courts can strike it down
- If individual assessment is unsupported by evidence, AAT can set it aside
- No statutory immunity from judicial review

Rationale: Immunity from invalidity creates legal black hole. If determinations cannot be invalidated, there is no practical remedy for unlawfulness. This is incompatible with rule of law.

Constitutional concern: Chapter III courts must have jurisdiction to adjudicate rights and obligations. Immunity from invalidity may violate separation of powers.

Schedule 2 failure: Section 92G(3) provides precisely this immunity - assessments remain valid even if unlawfully made.

Principle 6: Parliamentary Definition of Scope

Parliament, not Ministers, must define what can be assessed algorithmically.

- Primary legislation specifies what security assessments can address (terrorism, espionage, foreign interference, serious organized crime)
- Ministers cannot expand scope by regulation
- New risk categories require parliamentary amendment
- Parliament retains control over what is considered "security threat"

Rationale: Defining security threats is fundamental policy decision requiring parliamentary deliberation and approval. Delegating this to ministerial discretion enables scope creep without democratic accountability.

Schedule 2 failure: Section 92F(2) allows Minister to specify by legislative instrument (disallowable but not requiring affirmative parliamentary approval) what can be assessed.

Principle 7: Prohibition on Retrospective Operation

Assessments take effect when made, not retroactively.

- Algorithmic assessment applies from date determination actually taken
- Cannot retroactively substitute assessment made today for assessment made years ago
- Cannot retroactively void citizenship/visas based on re-assessment of old decisions
- Protections for legitimate expectations and settled matters

Rationale: Retrospective operation creates legal uncertainty and unfairness. People must be able to rely on determinations made at the time. Retrospective invalidation is hallmark of arbitrary rule.

Schedule 2 failure: Sections 92E(6)-(8) allow retrospective substitution of algorithmic assessments for previous human determinations.

Principle 8: Government Development and Operation Only

No private contractors for intelligence assessments.

- Algorithms developed by Australian government agencies
- Operated by Australian government personnel with security clearances
- No foreign corporations involved in core assessment process
- Private sector may provide technical infrastructure (servers, databases) but not assessment logic
- Source code and algorithms remain Commonwealth property
- Full audit access for IGIS and Auditor-General

Rationale: Intelligence assessment is core sovereign function. Foreign contractors create national security risk (what data extracted? what backdoors exist? who controls algorithms?). Private sector lacks accountability mechanisms applicable to government agencies.

Schedule 2 concern: Section 92E(1) "arrange for" authority enables Palantir (US corporation) involvement in intelligence assessment of Australian citizens/residents.

II. Illustrative Legislative Model (Non-Prescriptive)

This submission respectfully offers the Commission a specific legislative model demonstrating how these principles could be implemented. This is illustrative - the Commission may prefer different approaches - but shows concrete feasibility.

A. Proposed Legislative Structure

Intelligence Assessment (Algorithmic Decision-Support) Act 2026

This would be standalone legislation (not buried in Schedule to unrelated Bill) providing comprehensive framework for algorithmic intelligence assessment across all agencies (ASIO, AFP, AUSTRAC, Border Force, Defence Intelligence, etc.).

B. Key Provisions (Illustrative)

Part 1: Preliminary

Section 3: Objects

The objects of this Act are:

(a) To enable Australian intelligence and law enforcement agencies to use algorithmic decision-support tools to identify potential security threats efficiently; and

(b) To ensure such tools operate lawfully, transparently, and with appropriate safeguards for individual rights; and

(c) To maintain human decision-making for all determinations affecting individual rights, liberty, or status; and

(d) To provide independent review and accountability mechanisms; and

(e) To implement Recommendation 17.1 of the Royal Commission into the Robodebt Scheme regarding automated decision-making affecting Commonwealth administrative decisions.

Section 4: Definitions

algorithmic assessment: risk assessment or security evaluation generated wholly or partly by computer program processing data according to predetermined criteria

decision-maker: human officer authorized under this Act to make security assessments

methodology: the algorithms, data sources, weighting criteria, and decision rules used in algorithmic assessments (distinct from specific intelligence sources)

affected person: individual subject to adverse security assessment

Part 2: Algorithmic Assessment Framework

Section 10: Authorized Use

(1) An agency may use algorithmic decision-support for:
    (a) Terrorism risk assessment
    (b) Espionage risk assessment
    (c) Foreign interference risk assessment
    (d) Serious organized crime risk assessment
    (e) Sabotage risk assessment

(2) The Minister may not expand this list by legislative instrument.

(3) Parliament may expand this list by amending this section.

Section 11: Human Decision Requirement

(1) An algorithmic assessment does not constitute a security assessment for purposes of this or any other Act.

(2) A security assessment must be made by a human decision-maker after reviewing:
    (a) The algorithmic assessment output; and
    (b) All relevant circumstances known to the decision-maker; and

(c) Any response provided by the affected person; and
(d) Independent analysis where appropriate.

(3) The decision-maker must record in writing:
(a) Whether they accept, modify, or reject the algorithmic assessment; and
(b) The reasons for their determination; and
(c) What weight they gave to the algorithmic output versus other factors.

Section 12: Development and Operation

(1) Algorithmic assessment tools must be:
(a) Developed by Commonwealth agencies; or
(b) Developed by contractors under Commonwealth control with source code becoming Commonwealth property.

(2) Algorithmic assessment tools must be operated by:
(a) Australian government personnel with appropriate security clearances.

(3) No foreign corporation may:
(a) Operate algorithmic assessment tools; or
(b) Have access to assessment data or outputs; or
(c) Control or modify assessment algorithms without Commonwealth approval.

(4) The Inspector-General of Intelligence and Security and Auditor-General must have
full access to all algorithms, source code, training data, and outputs for audit
purposes.

Part 3: Transparency and Accountability

Section 20: Methodology Publication

(1) The Minister must table in Parliament within 6 months of commencement, and annually thereafter, a Methodology Statement describing:

(a) What data sources are considered in algorithmic assessments;
(b) What factors are weighted and how;
(c) How risk scores are calculated;
(d) What thresholds trigger adverse assessments;
(e) How often methodology is reviewed and updated;
(f) What validation and testing has been conducted;
(g) Whether any bias or discrimination has been identified and remediated.

(2) The Methodology Statement must be prepared in consultation with:
(a) The Inspector-General of Intelligence and Security;
(b) The Australian Human Rights Commission;
(c) Independent technical experts appointed by the Minister.

(3) The Methodology Statement must not disclose:
(a) Specific intelligence sources or methods;
(b) Identities of human sources;
(c) Ongoing investigations;
(d) Technical collection capabilities;
(e) Information that would enable evasion of detection.

(4) The Parliamentary Joint Committee on Intelligence and Security must review each
Methodology Statement and report to Parliament on whether it provides adequate

transparency while protecting legitimate security interests.

Section 21: Notification Requirement

(1) Before making an adverse security assessment that would affect an individual's:
    (a) Visa or citizenship status;
    (b) Employment or licensing;
    (c) Freedom of movement;
    (d) Access to government services;

    The decision-maker must provide the affected person with:

    (i) Notice that an algorithmic assessment has been generated;
    (ii) The risk category and score (if not compromising sources);
    (iii) Opportunity to respond within reasonable time;
    (iv) Information about review mechanisms.

(2) The notice must not disclose specific intelligence sources but must provide sufficient information for the person to meaningfully respond.

(3) Example: "You have been assessed as presenting elevated risk based on travel patterns, communications analysis, and financial transactions. You may provide
    explanation or contradictory evidence within 28 days."

(4) If disclosure would compromise ongoing investigation or imminent threat:
    (a) Notification may be delayed for up to 90 days;
    (b) Extension requires approval by senior officer and written justification;
    (c) Affected person must be notified as soon as operationally feasible.

Part 4: Independent Review

Section 30: Administrative Appeals Tribunal Jurisdiction

(1) An affected person may apply to the Administrative Appeals Tribunal for review
    of an adverse security assessment on the merits.

(2) The AAT may:
    (a) Affirm the assessment;
    (b) Vary the assessment;
    (c) Set aside the assessment and substitute its own;
    (d) Remit to the decision-maker with directions.

(3) The AAT must conduct review in accordance with security division procedures including:
    (a) Security-cleared AAT members;
    (b) In camera hearings where necessary;
    (c) Access to classified material;
    (d) Special advocates to represent affected person's interests when person cannot access classified material.

(4) The AAT may require the decision-maker to:
    (a) Disclose the algorithmic methodology;
    (b) Explain how algorithm reached its output;
    (c) Provide the specific data inputs that generated the assessment;
    (d) Justify the weighting and thresholds used.

(5) The AAT may receive confidential intelligence briefings from agencies but must:
    (a) Provide gist of information to affected person where possible;
    (b) Explain why information cannot be disclosed in full;
    (c) Consider appointing special advocate to test classified evidence.

Section 31: Inspector-General Oversight

(1) The Inspector-General of Intelligence and Security must:

    (a) Audit algorithmic assessment systems at least annually;
    (b) Review sample of assessments for lawfulness and reasonableness;
    (c) Investigate complaints about algorithmic assessments;
    (d) Report to Parliament on compliance with this Act;
    (e) Recommend methodology improvements;
    (f) Identify and report bias, discrimination, or systemic errors.

(2) IGIS must have full access to:
    (a) All algorithms and source code;
    (b) All training data and outputs;
    (c) All assessments made;
    (d) All decision-maker records;
    (e) All technical documentation.

(3) Agencies must not invoke security classification to refuse IGIS access.

(4) IGIS may require independent technical audit by external experts with appropriate clearances.

Part 5: Safeguards and Limitations

Section 40: No Immunity from Invalidity

(1) Nothing in this Act prevents a court from finding that an assessment is invalid
    if made unlawfully.

(2) An assessment found to be invalid is of no effect from the time it was made.

(3) Any decision made in reliance on invalid assessment is also of no effect unless
    court determines otherwise on equitable grounds.

(4) For avoidance of doubt: assessments are subject to judicial review under the Administrative Decisions (Judicial Review) Act 1977 and under section 75(v) of the Constitution.

Section 41: No Retrospective Operation

(1) An assessment takes effect when made, not retroactively.

(2) An assessment may not be substituted for a previous assessment to retroactively
    affect:
    (a) Citizenship or visa already granted;
    (b) Employment or licensing already obtained;
    (c) Any right or status already established.

(3) A new assessment may affect future applications or renewals but does not retroactively invalidate past decisions.

Section 42: Anti-Discrimination Requirements

(1) Algorithmic assessment methodology must not:
    (a) Discriminate on basis of race, ethnicity, religion, or national origin;
    (b) Use proxies for protected characteristics;
    (c) Perpetuate historical bias or discrimination.

(2) Before deploying algorithmic assessment, agency must:
    (a) Conduct bias testing across demographic groups;
    (b) Document disparate impact analysis;
    (c) Implement mitigation measures if bias detected;
    (d) Report findings to Australian Human Rights Commission.

(3) The Australian Human Rights Commission may:
    (a) Audit algorithmic systems for discrimination;
    (b) Require remediation;
    (c) Recommend suspension of discriminatory systems;
    (d) Report to Parliament on compliance with anti-discrimination law.

Section 43: Data Quality and Accuracy

(1) Agencies using algorithmic assessment must:
    (a) Maintain accurate, up-to-date data;
    (b) Verify data quality regularly;
    (c) Correct errors promptly when identified;
    (d) Notify affected persons if assessment based on inaccurate data.

(2) Affected person may request:
    (a) Access to data used in their assessment (subject to security protections);
    (b) Correction of inaccurate data;
    (c) Re-assessment if significant data errors identified.

(3) Decision-makers must consider data quality and reliability when weighting algorithmic output.

Part 6: Parliamentary Oversight

Section 50: Parliamentary Joint Committee Review

(1) The Parliamentary Joint Committee on Intelligence and Security must:

    (a) Review operation of this Act annually;
    (b) Review all Methodology Statements;
    (c) Receive briefings from IGIS on compliance;
    (d) Receive statistics on:
        (i) Number of assessments conducted;
        (ii) Number affirmed/varied/overturned on review;
        (iii) Demographic breakdown;
        (iv) Accuracy rates;
        (v) Complaints received;
    (e) Report to Parliament with recommendations for improvement.

(2) The PJCIS may conduct inquiries into specific concerns including:
    (a) Whether methodology is discriminatory;
    (b) Whether safeguards are adequate;
    (c) Whether agencies complying with this Act;
    (d) Whether amendments needed.

III. TRANSPARENCY ARCHITECTURE: BALANCING SECURITY AND ACCOUNTABILITY

This submission respectfully suggests that the common claim "we can't be transparent about intelligence methods" is often false. What actually needs protection is limited and specific. Most methodology can and should be public.

A. The False Dichotomy

False claim: "If we reveal how algorithmic assessment works, adversaries will evade detection."

Reality: Methodology transparency ≠ revealing specific sources or capabilities.

Example - What CAN be public:

> "Our algorithmic assessment considers: (1) Travel to conflict zones weighted by recency and frequency; (2) Communications with listed terrorist organizations weighted by nature and duration; (3) Financial transactions with designated entities weighted by amounts and purposes; (4) Associations with known threat actors weighted by proximity and frequency. These factors are combined into a risk score using weighted algorithm. Scores above threshold X trigger human review."

This reveals **what** is assessed and **how** weighting works. This is insufficient for evasion because:

- Adversaries still don't know **which specific communications** were detected
- Don't know **which financial transactions** were flagged
- Don't know **what technical collection capabilities** exist
- Don't know **which human sources** provided intelligence

Example - What MUST remain classified:

> "On [date], Five Eyes partner reported that [individual] met with [terrorist operative] in [location]. NSA intercept captured communication about [planned attack]. Financial intelligence shows transfer of $[amount] to [designated entity]."

This reveals **specific sources and capabilities** that must be protected.

B. International Models Demonstrating Feasibility

United Kingdom - MI5 Threat Assessment:

MI5 publicly describes methodology for assessing terrorism threats:

- What factors considered (capability, intent, vulnerability)
- How threats are graded (critical, severe, substantial, moderate, low)
- What triggers each level
- How often reviewed

MI5 does **not** disclose specific intelligence reports or collection capabilities.

Result: Public understands how threat assessment works. Courts can review methodology. Adversaries don't gain meaningful evasion capability.

Canada - CSIS Risk Assessment:

Canadian Security Intelligence Service provides Federal Court with:

- Complete description of algorithmic methodology
- Data sources used (categories, not specific reports)
- How risk scores calculated
- Validation testing results
- Bias analysis

Federal Court reviews in camera with security-cleared judges. Affected person receives summary (gist) enabling meaningful response without compromising sources.

Result: Judicial oversight robust. Discrimination claims can be assessed. Sources protected.

New Zealand - NZSIS Warrant Process:

NZ Security Intelligence Service must obtain judicial warrant for intrusive surveillance. Applications describe:

- Methodology for identifying targets
- What algorithmic tools assist human analysts
- How privacy is protected
- What safeguards prevent abuse

Commissioner of Intelligence Warrants reviews methodology, can require modifications, reports to Parliament on compliance.

Result: Algorithmic tools used lawfully under judicial supervision. Transparency and security coexist.

C. Australian Precedent - Existing Automated Systems with Transparency

Australia already operates automated systems in sensitive contexts with appropriate transparency:

ATO Compliance Risk Engine:

- Publicly documented methodology for tax compliance risk assessment
- Taxpayers know what triggers audit risk
- AAT reviews methodology for reasonableness
- No compromise to tax enforcement effectiveness

Border Force Advanced Passenger Processing:

- Publicly described methodology for identifying high-risk travelers
- Travelers know they may be subject to additional screening
- Ombudsman reviews for discrimination
- Effective border security maintained

Centrelink Payment Integrity:

- (Before Robodebt) Centrelink published criteria for debt investigation triggers
- Recipients knew what might cause review

- AAT reviewed methodology
- Fraud detection effective

These demonstrate: Transparency about methodology does not undermine operational effectiveness. It enables accountability without compromising legitimate objectives.

D. Proposed Transparency Levels for Intelligence Assessment

This submission suggests three tiers of transparency:

Tier 1 - Public (Methodology Statement tabled in Parliament):

- Categories of data sources (travel records, communications analysis, financial intelligence, human intelligence, open source)
- Factors weighted (what behaviors/associations/patterns considered)
- How risk scores calculated (weighted algorithm, threshold levels)
- Validation and accuracy testing results
- Bias and discrimination analysis
- How often methodology reviewed and updated
- How complaints investigated
- Statistical overview (numbers assessed, demographic breakdown, review outcomes)

Tier 2 - Parliamentary Committee (PJCIS in camera briefing):

- More detailed algorithmic specifications
- Specific agencies providing data
- Threshold score values
- Sample assessments (redacted)
- Technical validation details
- Comparison with international practices
- Identified vulnerabilities and remediation

Tier 3 - Inspector-General and Courts (full access):

- Complete source code
- All training data
- Individual assessment records
- Specific intelligence sources (in judicial review)
- Technical collection capabilities
- Ongoing investigations

This architecture enables:

- Public understanding of what government is doing
- Parliamentary oversight of methodology
- Judicial review of individual cases
- IGIS audit for compliance
- Protection of genuine secrets (sources, capabilities, investigations)

IV. ACCOUNTABILITY MECHANISMS: MAKING OVERSIGHT REAL

This submission respectfully suggests that accountability requires multiple overlapping mechanisms, not single point of failure.

A. Multi-Layered Oversight Model

Layer 1 - Internal Compliance (Agency Level):

- Designated Privacy Officers reviewing assessments for compliance
- Internal audit of algorithmic accuracy
- Regular bias testing
- Data quality assurance
- Decision-maker training and certification
- Complaint handling and correction mechanisms

Layer 2 - Inspector-General (Independent Statutory Oversight):

- IGIS for intelligence agencies (ASIO, ASIS, ASD, AGO, DIO, ONI)
- Commonwealth Ombudsman for law enforcement (AFP, Border Force)
- Full access to systems, data, and decision-making
- Annual audits with public reporting (classified annexes for sensitive details)
- Investigation of complaints
- Power to require remediation
- Reporting to Parliament

Layer 3 - Administrative Review (AAT):

- Merits review of individual assessments
- Security division procedures with cleared members
- In camera hearings protecting classified material
- Special advocates for affected persons
- Power to vary or set aside assessments
- Decisions published (with redactions for security)

Layer 4 - Judicial Review (Federal Court):

- Review of AAT decisions on questions of law
- Constitutional challenges to methodology
- Discrimination claims under Racial Discrimination Act
- Human rights claims under International Covenant
- Full judicial process with protections for classified material

Layer 5 - Parliamentary Oversight (PJCIS):

- Annual review of methodology statements
- Inquiry power into systemic concerns
- Briefings from IGIS, agencies, affected persons
- Recommendations to Parliament

- Public hearings (where possible) and classified briefings (where necessary)
- Report to Parliament with recommendations for legislative amendment

Layer 6 - Human Rights Commission:

- Audit algorithmic systems for discrimination
- Receive complaints about bias
- Investigate systemic discrimination
- Recommend remediation
- Report to Parliament on human rights compliance
- Intervene in AAT/Court proceedings on discrimination issues

Layer 7 - Auditor-General:

- Audit expenditure on algorithmic systems
- Performance audit of accuracy and effectiveness
- Value-for-money analysis
- Review contracting arrangements
- Report to Parliament on financial accountability

No single layer is sufficient. Together, they create robust accountability:

- Agencies have incentive to comply (IGIS will audit)
- Individuals can challenge unfair assessments (AAT review)
- Systemic problems surface through multiple channels (IGIS, AAT patterns, PJCIS inquiry)
- Discrimination identified and remediated (AHRC oversight)
- Parliament informed and can amend (PJCIS recommendations)
- Courts enforce constitutional limits (judicial review)
- Public expenditure accountable (Auditor-General)

B. Special Advocate Model for Classified Material

A persistent challenge: How to provide fair hearing when some evidence must remain classified?

This submission suggests adopting the **Special Advocate model** successfully used in UK, Canada, and New Zealand:

How it works:

1. Affected person applies to AAT for review of adverse assessment
2. AAT determines some evidence must remain classified for security reasons
3. AAT appoints Special Advocate - security-cleared lawyer representing affected person's interests
4. Special Advocate receives full classified brief, can question witnesses, challenge evidence
5. Special Advocate cannot disclose classified details to affected person but can:
    - Challenge whether classification justified
    - Test reliability of classified evidence
    - Argue for maximum disclosure consistent with security
    - Provide gist/summary to affected person where possible
    - Ensure affected person's perspective considered even in closed proceedings

Benefits:

- Affected person has advocate even in closed proceedings
- Classified evidence is tested, not accepted uncritically
- Agencies can't hide weak evidence behind classification
- Fairness improved while protecting genuine secrets

Australian precedent: Special advocates already used in:

- Control order proceedings (Criminal Code)
- Preventative detention orders (Criminal Code)
- Certain national security visa cancellations

This submission suggests: Special advocates should be standard in all algorithmic assessment reviews involving classified material.

C. Bias Testing and Remediation Requirements

Algorithmic systems can perpetuate or amplify existing discrimination. This submission suggests mandatory bias testing:

Pre-Deployment Testing:
Before any algorithmic assessment system is deployed:

1. Test against representative sample across demographic groups
2. Measure disparate impact (are some groups assessed as higher risk at higher rates?)
3. Investigate causes of disparity (is it justified by evidence or proxy discrimination?)
4. Implement mitigation if unjustified disparity found
5. Document all testing and submit to AHRC for review
6. AHRC approval required before deployment

Ongoing Monitoring:
After deployment:

1. Quarterly analysis of assessment outcomes by demographic groups
2. Investigation of any emerging disparities
3. Re-calibration of algorithms if bias detected
4. Reporting to AHRC and PJCIS
5. Public reporting (aggregated, anonymized statistics)

Independent Audit:
Annual independent technical audit by external experts:

1. Review training data for bias
2. Analyze algorithmic logic for discriminatory patterns
3. Test against international best practices
4. Recommend improvements
5. Report to IGIS, AHRC, and PJCIS

Remediation Process:
If bias identified:

1. Immediate suspension of discriminatory component
2. Analysis of harm caused
3. Notification to affected persons
4. Opportunity for re-assessment
5. Compensation where appropriate
6. Implementation of bias-free alternative
7. Public reporting on remediation

This demonstrates: Preventing algorithmic discrimination is technically feasible and legally required.

V. IMPLEMENTATION PATHWAY: FROM HERE TO THERE

This submission respectfully offers the Commission a phased implementation pathway showing how to transition from current flawed Schedule 2 framework to legitimate algorithmic governance:

Phase 1 - Immediate (0-6 months): Stop the Harm

Actions:

1. Attorney-General advises non-operationalization of sections 92E-92G pending review
2. No algorithmic assessments conducted under Schedule 2 authority
3. No contracts signed with private sector (including Palantir) for assessment tools
4. Parliamentary inquiry established to reconsider framework
5. Solicitor-General constitutional advice obtained
6. IGIS investigation into ASIO's role in Schedule 2 development
7. Drafting begins on replacement legislation

Rationale: Prevent harm while replacement framework developed. No one harmed by delay because system not yet operational.

Phase 2 - Short Term (6-12 months): Develop Replacement Framework

Actions:

1. Parliamentary inquiry receives submissions from:
    - Civil liberties organizations
    - Legal experts
    - Technical experts
    - Affected communities
    - International comparators
    - Intelligence and law enforcement agencies
2. Draft legislation developed incorporating:
    - Eight core principles outlined above

- Multi-layered accountability mechanisms
- Transparency requirements
- Anti-discrimination safeguards
- Human decision-making requirement

3. Exposure draft released for public consultation
4. Parliamentary Joint Committee on Intelligence and Security reviews draft
5. Government introduces bill with full explanatory materials including:
    - Robodebt Royal Commission Recommendation 17.1 compliance
    - Constitutional advice
    - International comparison
    - Human rights compatibility
6. Adequate parliamentary scrutiny period (minimum 90 days)

Rationale: Democratic legitimacy requires informed parliamentary consent after genuine deliberation, consistent with international standards for automated decision-making frameworks affecting individual rights.

Phase 3 - Medium Term (12-24 months): Build Capability

Actions:

1. Replacement legislation passed
2. Agencies begin developing algorithmic tools under new framework:
    - Government-developed or government-controlled
    - No foreign contractor involvement in core assessment
    - Full audit access for IGIS
    - Commonwealth ownership of intellectual property
3. Methodology Statement drafted in consultation with IGIS, AHRC, independent experts
4. Bias testing conducted on pilot systems
5. Special Advocate panel established and security-cleared
6. AAT security division capacity expanded
7. PJCIS receives first briefing on methodology
8. Public education campaign explaining how system works and what rights people have

Rationale: Building legitimate capability takes time but prevents Robodebt-style rush-to-deployment resulting in unlawful system.

Phase 4 - Long Term (2-5 years): Full Operation with Oversight

Actions:

1. Algorithmic assessment tools operational under new framework
2. All eight core principles implemented:

    - Human decision-making ✓
    - Notification before adverse action ✓
    - Methodology transparency ✓
    - Independent merits review ✓
    - No immunity from invalidity ✓
    - Parliamentary scope control ✓
    - No retrospective operation ✓
    - Government-only operation ✓
3. Multi-layered oversight functioning:
    - IGIS annual audits ✓
    - AAT reviewing individual cases ✓
    - PJCIS parliamentary oversight ✓
    - AHRC bias monitoring ✓
    - Auditor-General financial audit ✓
4. Annual reporting to Parliament on:
    - Number of assessments
    - Demographic breakdown
    - Accuracy rates
    - AAT outcomes
    - Bias testing results
    - Complaints and remediation
5. International engagement sharing Australian model
6. Regular reviews and amendments as technology and threats evolve

Rationale: Legitimate algorithmic governance is ongoing process, not one-time implementation. Continuous oversight and adaptation required.

VI. INTERNATIONAL COMPARISON: LEARNING FROM OTHERS

This submission respectfully suggests the Commission can recommend Australia learn from democracies successfully using algorithmic tools with proper safeguards:

A. European Union - AI Act (2024)

Framework:

- Categorizes AI systems by risk level
- "High risk" systems (law enforcement, immigration) face strict requirements:
    - Human oversight mandatory
    - Transparency about how system works
    - Data quality standards
    - Bias testing required

    - Right to explanation
    - Independent conformity assessment

Enforcement:

- Significant fines for non-compliance (up to 6% global revenue)
- National competent authorities audit systems
- Individuals can complain to authorities
- Courts can order remediation or suspension

Lesson for Australia: Comprehensive regulatory framework prevents deployment of opaque or discriminatory systems.

B. United Kingdom - Algorithmic Transparency Standard

Framework:

- Government departments using algorithmic tools must publish:
    - What tool does
    - How it works
    - What data it uses
    - How decisions are made
    - How to challenge decisions

Oversight:

- Information Commissioner audits for data protection compliance
- Equality and Human Rights Commission audits for discrimination
- Courts review under judicial review jurisdiction
- Parliamentary committees inquire into specific concerns

Notable example - MI5 threat assessment:

- Methodology public while sources protected
- Independent Investigatory Powers Commissioner reviews
- Courts can assess in closed proceedings
- Has not compromised operational effectiveness

Lesson for Australia: Transparency about methodology compatible with protecting sources and operational security.

C. Canada - Algorithmic Impact Assessment

Framework:

- All federal agencies using automated decision systems must complete Algorithmic Impact Assessment (AIA)
- AIA evaluates:
    - Impact level (low, medium, high, very high)
    - Risk of bias or discrimination
    - Transparency appropriate to impact level

    - Human oversight required
    - Review mechanisms

Oversight:

- Treasury Board Secretariat reviews all AIAs
- Privacy Commissioner audits data protection
- Canadian Human Rights Commission audits discrimination
- Federal Court reviews under judicial review
- Affected persons can request explanation and review

Notable example - Immigration risk assessment:

- CBSA uses algorithmic tools for traveler risk assessment
- Methodology reviewed by Privacy Commissioner
- Found some bias issues, required remediation
- Now operates with better safeguards

Lesson for Australia: Pre-deployment impact assessment identifies risks before harm occurs.

D. New Zealand - Algorithm Charter

Framework:

- Government agencies using algorithms commit to:
    - Transparency about use
    - Partnership with affected communities
    - Human oversight
    - Embedding te Tiriti o Waitangi (Treaty principles)
    - Bias mitigation

Oversight:

- Privacy Commissioner audits
- Chief Ombudsman reviews complaints
- Statistics NZ monitors demographic impacts
- Public reporting on algorithm use

Notable example - MSD fraud detection:

- After concerns about bias, methodology reviewed
- Found disparate impact on Māori and Pacific communities
- Required remediation and oversight
- Now operates with better fairness safeguards

Lesson for Australia: Community partnership and cultural sensitivity improve both fairness and effectiveness.

E. What Australia Can Learn

Common features of successful international models:

1. Transparency about methodology (not sources)
2. Human oversight required for high-impact decisions
3. Mandatory bias testing before deployment
4. Independent oversight with real powers
5. Right to explanation and review
6. Public reporting on use and outcomes
7. Remediation when problems identified
8. No immunity from legal review

What Australia should avoid (lessons from failures):

1. Opaque systems without explanation (Netherlands - welfare fraud detection suspended after discrimination findings)
2. Automation without human review (UK - A-level exam algorithm disaster)
3. Deployment without bias testing (US - COMPAS recidivism tool found racially biased)
4. Private contractor control (multiple jurisdictions - vendor lock-in and accountability gaps)

This submission suggests: Australia can be international leader by implementing comprehensive framework learning from both successes and failures elsewhere.

VII. TECHNICAL FEASIBILITY: ADDRESSING COMMON OBJECTIONS

This submission respectfully addresses common objections to transparency and accountability in algorithmic intelligence assessment:

Objection 1: "Transparency enables evasion"

Response:
Methodology transparency ≠ revealing specific sources or capabilities.

Adversaries already assume surveillance exists. Knowing that government monitors "travel to conflict zones, communications with listed entities, financial transactions with designated groups" does not enable meaningful evasion because:

- They don't know which specific travels/communications/transactions were detected
- They don't know what technical collection capabilities exist
- They don't know thresholds or weighting
- They don't know when they're being monitored

Example: Publishing tax compliance criteria doesn't enable tax evasion because knowing "large cash transactions trigger audit risk" doesn't tell you whether your specific transaction will be detected.

Objection 2: "Courts can't review classified intelligence"

Response:
Courts routinely review classified material in:

- Control order proceedings
- Preventative detention orders
- National security visa cancellations
- Closed evidence applications
- Security clearance disputes

Procedures exist:

- Security-cleared judges
- In camera hearings
- Special advocates
- Gist disclosure to affected persons
- Protective orders

This is routine judicial practice, not novel or impractical.

Objection 3: "Human review slows down process"

Response:
Algorithms provide speed in identifying cases requiring human attention. Humans then review efficiently:

- Algorithm processes 10,000 travelers → flags 50 requiring review → humans assess 50 (not 10,000)
- This is faster than human review of all 10,000
- This is slower than automated rejection of all 50
- **But**: Speed does not justify unlawful automation

Robodebt lesson: Speed (processing debts quickly) does not justify illegality (income averaging unlawful).

Objection 4: "Government can't afford extensive oversight"

Response:
Oversight is less expensive than Royal Commissions after harm.

Robodebt cost:

- $1.8 billion in unlawful debts
- $751 million in compensation
- Royal Commission cost
- Incalculable reputational harm

Prevention cost (if proper oversight existed):

- AAT reviews: ~$5,000 per case

- IGIS auditing: ~$2 million annually
- Bias testing: ~$500,000 per system
- Special advocates: ~$200,000 annually

Preventing one Robodebt saves hundreds of millions and immense human suffering.

This submission suggests: "We can't afford oversight" is false economy.

Objection 5: "Private sector more capable than government"

Response:
Technical capability can be acquired through:

- Government development (ATO, Border Force already do this successfully)
- Contractors developing under government control with IP transfer
- Government hiring technical expertise
- Academic partnerships for algorithm development
- International collaboration (Five Eyes sharing non-commercial tools)

What government should not do:

- Outsource core sovereign function (intelligence assessment) to foreign corporations
- Create dependency on proprietary systems without source code access
- Enable foreign access to Australian intelligence data
- Lose capability to audit or modify systems

Schedule 2 concern: Palantir involvement creates all these risks.

Objection 6: "Anti-discrimination testing is impractical"

Response:
Testing is standard practice in:

- Medical device approval (FDA testing requirements)
- Financial services (lending algorithm bias testing)
- Employment systems (EEOC compliance testing)
- Education admissions (disparate impact analysis)

Process:

1. Test algorithm against representative sample
2. Measure outcomes across demographic groups
3. If disparate impact found, investigate whether justified by legitimate factors
4. If not justified, modify algorithm or abandon
5. Document all testing
6. Subject to independent audit

Cost: ~$500,000 per system initially, ~$100,000 annually for ongoing monitoring.

Benefit: Prevents discriminatory system causing immense harm requiring expensive remediation.

This submission suggests: Bias testing is technically feasible, legally required, and cost-effective prevention.

VIII. CONCLUSION: TECHNOLOGY CAN SERVE JUSTICE

This Part has demonstrated that legitimate algorithmic intelligence assessment is achievable.

The choice is not:

- Option A: Schedule 2 (immunity, opacity, automation)
- Option B: No technology (manual processing only)

The choice is:

- Option A: Schedule 2 (unlawful automation enabling Robodebt repetition)
- Option B: Legitimate framework (technology serving justice with accountability)

This submission has shown:

✓ **Eight core principles** are technically feasible and legally sound

✓ **Detailed legislative framework** demonstrating concrete implementation

✓ **Transparency architecture** showing methodology disclosure compatible with source protection

✓ **Multi-layered accountability** creating robust oversight without single point of failure

✓ **Implementation pathway** from current state to legitimate system in phased approach

✓ **International precedents** proving democracies successfully use algorithms with proper safeguards

✓ **Technical feasibility** addressing and refuting common objections

What this requires:

🔹 Political will to prioritize lawfulness over speed

🔹 Parliamentary scrutiny with adequate time and information

🔹 Investment in government capability (not outsourcing to foreign contractors)

🔹 Acceptance that security and accountability are compatible, not contradictory

🔹 Learning from Robodebt rather than repeating it

🔹 Recognizing that immunity from legal review is incompatible with rule of law

What this achieves:

✅ Effective intelligence assessment (technology assists human decision-makers)

✅ Protection of rights (notification, review, no immunity)

✅ Democratic accountability (Parliament controls scope, PJCIS oversees)

✅ Non-discrimination (bias testing and remediation required)

✅ Compliance with Recommendation 17.1 (clear review path, departmental explanation, algorithm availability)

✅ Constitutional soundness (courts can review, no immunity from Chapter III jurisdiction)

✅ International leadership (Australia as model for algorithmic governance done right)

This submission respectfully urges the Commission to recommend:

Parliament should repeal Schedule 2 and replace it with a comprehensive framework implementing the principles outlined in this Part. Detailed legislative language can be developed through proper parliamentary process with adequate consultation and scrutiny.

The Commission can show Parliament: Rejecting Schedule 2 is not rejecting technology. It is requiring technology to serve justice, respect rights, and operate under law.

Robodebt Royal Commission Recommendation 17.1 showed the path. This submission has shown that path is followable.

The question is whether Parliament has the will to follow it.

APPENDICES

This submission includes the following appendices providing supporting documentation, timelines, and reference materials for the Commission's consideration.

APPENDIX A: TIMELINE OF KEY EVENTS

This timeline presents key events relevant to Schedule 2, Robodebt, and the broader pattern documented in this submission.

The Robodebt Era (2015-2019)

2015

- **July**: Department of Human Services begins income averaging for debt calculation
- No legislative authority for this practice
- No notification to recipients of methodology change

2016

- **January-December**: Robodebt program expands
- Automated debt notices sent without human review
- Recipients not informed income averaging is being used
- No clear review pathway

2017

- **September**: Secretary of Department begins raising concerns internally
- Legal advice sought on lawfulness of income averaging
- Advice indicates potential problems
- Program continues

2018

- **November**: Victoria Legal Aid challenges Robodebt lawfulness
- Settlement reached with Commonwealth but no public admission of unlawfulness
- Program continues despite legal challenge

2019

- **May**: Federal election - Coalition government returned
- Robodebt continues despite ongoing legal concerns
- **September**: Gordon Legal launches class action
- Commonwealth begins defending Robodebt as lawful
- **November**: Federal Court proceedings begin

Robodebt Ending and Royal Commission (2020-2023)

2020

- **May**: Commonwealth discontinues Robodebt defense
- Admits income averaging was unlawful
- Agrees to repay debts and compensate victims
- **November**: Settlement approved - $1.8 billion in unlawful debts, $751 million compensation

2021

- **August**: Tampa incident 20th anniversary
- Offshore detention continues with bipartisan support

2022

- **May**: Federal election - Labor government elected under Anthony Albanese
- **August**: Prime Minister announces Royal Commission into Robodebt Scheme
- **October**: Royal Commission begins hearings
- Terms of reference include examination of:
    - How unlawful scheme occurred
    - Who was responsible
    - What reforms needed to prevent repetition

2023

- **January-June**: Royal Commission hearings continue
- Evidence of concealment, legal advice ignored, ministerial knowledge
- Personal stories of harm and deaths
- **July 7**: Royal Commission Report delivered to Governor-General
- 56 recommendations including Recommendation 17.1 on automated decision-making
- Multiple referrals for civil/criminal investigation
- **September**: Royal Commission Report publicly released

2024

- **Throughout year**: Government considers Royal Commission recommendations
- Some implementation begins (NACC referrals, apology, compensation)
- Recommendation 17.1 on automated decision-making not yet implemented

Schedule 2 Development and Passage (2025-2026)

2025

- **March-June**: Schedule 2 development (timeframe estimated)
- ASIO and Home Affairs drafting provisions
- Consultation occurs but scope unknown
- Robodebt Royal Commission findings not integrated
- **July-September**: Cabinet consideration
- Cabinet submissions not public
- National Security Committee involvement likely
- **November**: Palantir Technologies receives Protected-level security clearance
- Timing: months before parliamentary introduction
- Clearance enables access to classified intelligence systems
- No public disclosure of clearance purpose or relationship to pending legislation
- **December**: Bill introduced to Parliament
- *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Bill 2026*
- 185-page Explanatory Memorandum
- No mention of Robodebt Royal Commission
- No mention of Recommendation 17.1
- No explanation of immunity from invalidity
- No disclosure of private sector deployment authority

2026

- **14 January**: Exposure draft released (9am)
- **15 January**: IGIS submits Submission 145 to PJCIS
- Warning about Schedule 2 automation risks
- Recommends enhanced safeguards and oversight mechanisms
- **16 January**: Submissions close (5pm)
- Total consultation period: 54 hours
- **17 January**: PJCIS hearings
- **20 January (morning)**: PJCIS report tabled
- Bill split from omnibus into two separate Acts
- State law incorporation introduced for first time
- **20 January (afternoon)**: Senate debate begins
- Senator McDonald mentions Robodebt connection once

- No other substantive engagement with automation provisions
- **21 January**: House of Representatives debate
- Focus on firearms and antisemitism
- Minimal discussion of Schedule 2 provisions
- **21 January**: Royal Assent
- *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* becomes law
- Sections 92E-92G now in force

The Broader Pattern Context

2001

- **August 26**: MV Tampa rescues 433 asylum seekers
- Australian government refuses entry
- Political crisis during election campaign
- **October 7**: Children Overboard claims begin
- Government claims asylum seekers threw children overboard
- Senate inquiry later confirms claims were false
- Narrative maintained through election campaign
- **November 10**: Federal election
- Howard Coalition government returned with increased majority
- Tampa and Children Overboard credited as significant electoral factors
- Offshore detention becomes entrenched bipartisan policy

2001-2008

- First iteration of offshore detention (Nauru, Manus Island)
- Thousands detained
- Significant costs and documented harms

2008

- Rudd Labor government closes offshore detention centres

2012

- Gillard Labor government reopens offshore detention
- Bipartisan support re-established

2012-present

- Continuous offshore detention operation (14 years)

- 4,183 people sent offshore
- ~110 still on Nauru (as of August 2025)
- ~70-80 in PNG
- 14 confirmed deaths
- $10+ billion expenditure
- No Royal Commission despite 24 years total operation

APPENDIX B: OFFSHORE DETENTION STATISTICS

This appendix provides comprehensive statistics on offshore detention relevant to the pattern documented in Part IIB.

Population Statistics

Total sent offshore (2012-present): 4,183 people

Current detention (as of August 2025):

- Nauru: ~110 people
- PNG: ~70-80 people
- Total: ~180-190 people

Status of those currently detained:

- 78 on Nauru have been recognized as refugees (entitled to protection under international law)
- Many detained for 10+ years without resolution

Historical numbers:

- Peak detention (2014): ~2,000 people offshore
- Gradual reduction through resettlement in third countries (primarily United States)
- ~3,800 resettled or returned (many after 5-10 years detention)

Duration Statistics

Length of detention:

- 10+ years: Dozens of people
- 5-10 years: Hundreds of people
- Average detention period: ~5 years before resettlement/return

Longest detentions:

- Some individuals detained 2012-present: 13+ years continuously
- Children born in detention, entire childhood in detention facilities
- Families separated (some members resettled, others remaining detained)

Deaths and Serious Harm

Confirmed deaths (2012-present): 14

Breakdown by cause:

- Suicides: 7 confirmed
    - Including 2 self-immolations
- Medical-related: 3 (preventable deaths from delayed treatment)
- Other causes: 4 (including violence, accidents)

Serious self-harm incidents:

- Nauru Files (2016 leak): 2,116 incident reports documenting:
    - Sexual assault incidents
    - Hundreds of self-harm incidents
    - Mental health crises
    - Physical abuse allegations
    - Attempted suicides

Medical assessments:

- Médecins Sans Frontières (2018): Offshore detention mental health suffering "among most severe we have ever witnessed" globally
- Australian Medical Association: Repeated calls for closure on medical grounds
- Royal Australasian College of Physicians: Children suffering "catastrophic" harm

Financial Statistics

Total expenditure (2001-2008 + 2012-present): Estimated $10+ billion

Cost breakdown:

- 2012-present (14 years): ~$7-8 billion documented
- 2001-2008 (7 years): ~$2-3 billion estimated

Per-person costs:

- Annual cost per person (Nauru): ~$3.4 million
- For current population (~110 on Nauru): ~$374 million annually

Comparison with onshore processing:

- Onshore processing cost: ~$90,000-100,000 per person annually
- Offshore processing cost: ~$3.4 million per person annually
- Offshore is approximately 37 times more expensive than onshore

Contractor Issues

Paladin affair:

- Contract value: $423 million (expanded to $532 million over 22 months)
- Company had minimal prior experience in refugee services

- Registered office was residential property on Kangaroo Island
- Contract awarded without competitive tender
- Parliamentary scrutiny faced delays and commercial confidentiality claims

Canstruct concerns:

- Contract escalation: $8 million (September 2017) to $385 million (October 2017)
- Total payments 2017-present: Approximately $1.5 billion
- Allegations emerged regarding director's immigration fraud charges
- Questions raised about vetting and oversight processes

Unauthorized expenditure (2017 audit):

- Total unauthorized spending identified: $2.2 billion
- Breakdown:
    - $1.1 billion approved by officers without proper delegation
    - $1.1 billion with no authorization record
- No officials disciplined, programs continued

International Law Considerations

United Nations Human Rights Committee (January 2025):

- Finding: Australia violated International Covenant on Civil and Political Rights
- Violations identified: Arbitrary detention, cruel/inhuman treatment
- Recommended remedy: Immediate release, compensation, policy change
- Australia's response: No compliance with recommendations

International Criminal Court Prosecutor (2020):

- Preliminary examination conducted
- Assessment: Conduct warranted investigation consideration
- Status: No formal investigation opened (jurisdictional considerations)

Papua New Guinea Supreme Court (2016):

- Finding: Detention on Manus Island unconstitutional under PNG law
- Ordered: Closure of detention centre
- Australia's response: Eventually closed Manus centre (2017), continued detention in alternative PNG locations

---

APPENDIX C: GLOSSARY OF TERMS

This appendix defines technical terms used throughout this submission.

AAT (Administrative Appeals Tribunal): Independent tribunal reviewing Commonwealth administrative decisions on the merits.

AFP (Australian Federal Police): National police service, investigates Commonwealth crimes.

Algorithmic assessment: Risk assessment or security evaluation generated wholly or partly by computer program processing data according to predetermined criteria.

ASIO (Australian Security Intelligence Organisation): Australia's domestic security intelligence agency, responsible for counter-terrorism, counter-espionage, counter-foreign interference.

Bipartisan: Supported by both major parties (Labor and Coalition), limiting effective parliamentary opposition.

Chapter III (Constitution): Establishes judicial power of Commonwealth, vests in High Court and federal courts.

Children Overboard: False 2001 narrative claiming asylum seekers threw children overboard to manipulate rescue. Senate inquiry confirmed falsehood.

EM (Explanatory Memorandum): Document accompanying bill explaining purpose and effect of each provision, intended to inform parliamentary debate.

IGIS (Inspector-General of Intelligence and Security): Independent statutory officer overseeing Australian intelligence agencies for legality, propriety, and compliance with human rights.

Immunity from invalidity: Legal provision making administrative decisions valid even if made unlawfully. Prevents courts from invalidating unlawful decisions.

NACC (National Anti-Corruption Commission): Federal integrity agency investigating corruption by Commonwealth public officials.

Offshore detention: Policy of detaining asylum seekers who arrive by boat in processing centres on Nauru and PNG rather than in Australia. Continuous operation 2012-present, previously 2001-2008.

Palantir Technologies Inc.: US corporation providing data analytics and algorithmic assessment tools, received Australian Protected-level security clearance November 2025.

PJCIS (Parliamentary Joint Committee on Intelligence and Security): Bipartisan committee overseeing intelligence agencies and reviewing national security legislation.

Proprietary algorithm: Computer program owned by private corporation, source code not publicly available, often protected as trade secret.

Recommendation 17.1: Robodebt Royal Commission recommendation requiring automated decision-making systems to provide: clear review pathways for affected persons, transparency about how algorithms operate, and independent scrutiny mechanisms.

Retrospective/retroactive operation: Law or decision applying to events that occurred before the law was passed or decision was made.

Robodebt: Unlawful automated debt scheme 2015-2019, calculated social security overpayments using income averaging (unlawful method), affected 443,000 people, $1.8 billion in unlawful debts.

Schedule 2: Sections 92E-92G of ASIO Act inserted by *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026*, authorizing algorithmic intelligence assessment.

Section 92E: Authorizes ASIO Director-General to arrange for computer programs to make specified assessment action.

Section 92F: Allows Minister to define "specified assessment action" via non-legislative instrument.

Section 92G: Provides that specified assessment actions are not invalid due to failure to comply with safeguard requirements (immunity from invalidity).

Section 36(1)(d): Previously required security assessment be given to person adversely affected; Schedule 2 deletes this requirement.

Tampa: Norwegian freighter MV Tampa that rescued 433 asylum seekers near Australia August 2001, Australian government refused entry triggering political crisis and policy shift to offshore detention.

APPENDIX D: REFERENCES AND SOURCES

This appendix lists key sources referenced throughout this submission.

Primary Legal Sources

Legislation:

- *Australian Security Intelligence Organisation Act 1979* (Cth)
- *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* (Cth)
- *Administrative Decisions (Judicial Review) Act 1977* (Cth)
- *Racial Discrimination Act 1975* (Cth)
- Australian Constitution

Royal Commission Reports:

- Royal Commission into the Robodebt Scheme, Report (July 2023)

Parliamentary Materials:

- Explanatory Memorandum, *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Bill 2026*
- House of Representatives Hansard, 20-21 January 2026
- Senate Hansard, 20-21 January 2026
- Parliamentary Joint Committee on Intelligence and Security, Submission 145 from Inspector-General of Intelligence and Security (15 January 2026)

Constitutional Analysis:

- Twomey, Anne. "The messy new laws to ban 'hate' groups" *The Conversation* (21 January 2026)
- Twomey, Anne. "Constitutional Clarion" YouTube channel analysis (21 January 2026)

Case Law:

- *Plaintiff S157/2002 v Commonwealth* (2003) 211 CLR 476 (privative clauses)
- *Chu Kheng Lim v Minister for Immigration* (1992) 176 CLR 1 (detention limits)
- *McCloy v New South Wales* [2015] HCA 34 (proportionality test)
- *Wertheim v Haddad* [2025] FCA 720 (Israel criticism not inherently antisemitic)

Government and Official Sources

Australian Government:

- Department of Home Affairs annual reports and statistics
- ASIO annual reports
- Inspector-General of Intelligence and Security annual reports
- Australian National Audit Office reports
- Auditor-General reports on offshore detention expenditure

International Bodies:

- United Nations Human Rights Committee decisions on Australia
- International Criminal Court Prosecutor statements
- PNG Supreme Court decisions on Manus detention
- UNHCR statements on offshore detention

Academic and Expert Sources

Books:

- Arendt, Hannah. *The Origins of Totalitarianism* (1951, updated editions)
- Arendt, Hannah. *Eichmann in Jerusalem: A Report on the Banality of Evil* (1963)
- Marr, David & Wilkinson, Marian. *Dark Victory* (2004) [Tampa and Children Overboard]

Academic Articles and Books on Algorithmic Governance:

- Pasquale, Frank. *The Black Box Society* (2015) [algorithmic accountability]
- Eubanks, Virginia. *Automating Inequality* (2018) [automated decision-making harms]
- O'Neil, Cathy. *Weapons of Math Destruction* (2016) [algorithmic bias]
- Yeung, Karen. "Algorithmic regulation: A critical interrogation" (2018) *Regulation & Governance*

Media and Investigative Sources

Offshore Detention:

- The Guardian Australia, "Nauru Files" investigation (2016)
- ABC Four Corners, investigations into offshore detention
- Médecins Sans Frontières, reports on mental health in offshore detention

Robodebt:

- ABC investigations into Robodebt (2017-2020)
- The Guardian Australia reporting

- Senate Committee inquiries

Palantir:

- Michael West Media reporting on Palantir clearance (19 January 2026)
- Capital Brief reporting on Palantir lobbying activities (July 2025)
- Austender contract disclosures

International Comparisons

European Union:

- EU Artificial Intelligence Act (2024)
- EU General Data Protection Regulation (GDPR)

United Kingdom:

- UK Algorithmic Transparency Standard
- Information Commissioner's Office guidance on AI
- Cabinet Office algorithmic impact assessment frameworks

Canada:

- Algorithmic Impact Assessment framework
- Treasury Board Directive on Automated Decision-Making

New Zealand:

- Algorithm Charter for Aotearoa New Zealand

NGO and Civil Society Sources

Human Rights Organizations:

- Australian Human Rights Commission reports
- Human Rights Watch reports on offshore detention
- Amnesty International reports

Legal Organizations:

- Law Council of Australia submissions
- Public Interest Advocacy Centre materials

Refugee Advocacy:

- Refugee Council of Australia reports and statistics
- UNHCR Australia statements

End of Appendices

These appendices provide supporting documentation for the submission's analysis and recommendations.

CONCLUSION

From Analysis to Action

This submission has examined Schedule 2 of the *Combatting Antisemitism, Hate and Extremism (Firearms and Customs Laws) Act 2026* through multiple analytical lenses: definitional frameworks, institutional design, historical patterns, parliamentary process, statutory construction, and operational risks. The analysis has been detailed and at times technical, reflecting the complexity of the subject matter and the seriousness with which this submission approaches both the problem of antisemitism and the proposed legislative response.

It is appropriate now to step back and state plainly what this submission has sought to establish, and what it respectfully asks the Commission to consider.

The Core Argument

Antisemitism is real, harmful, and deserving of serious response. Jewish Australians have experienced increased fear, vandalism, harassment, and violence. These concerns are legitimate and require action.

However, the question is not whether to respond, but how to respond effectively.

This submission has documented that Schedule 2:

- was enacted following a legislative process that raises questions about transparency and informed parliamentary consideration;
- replicates structural features that the Robodebt Royal Commission identified as problematic—automated decision-making, immunity provisions, limited notification, constrained review mechanisms—enacted within 28 months of that Commission's findings;
- operates within contested definitional frameworks that create uncertainty about enforcement boundaries;
- may enable applications beyond the firearms licensing context emphasized during parliamentary debate;
- introduces AI-specific operational risks including hallucination, feedback loops, efficiency paradoxes, and oversight challenges; and
- was developed without adequate incorporation of Recommendation 17.1's safeguards regarding automated systems affecting individuals.

These are not hypothetical concerns. They are grounded in:

- documented recent Australian experience (Robodebt);
- specific statutory language in Schedule 2;
- parliamentary records showing minimal substantive debate;
- expert warnings from IGIS that were issued but not incorporated;
- international evidence of how similar systems operate; and

- established constitutional and administrative law principles.

What This Submission Does Not Claim

In the interest of clarity, it is worth stating explicitly what this submission **does not** argue:

This submission does not claim:

- that antisemitism is not a serious problem requiring response;
- that Jewish Australians' safety concerns are illegitimate or exaggerated;
- that ASIO, Home Affairs, or drafting officials acted with improper intent;
- that Schedule 2 will inevitably be misused or that harm is predetermined;
- that contemporary Australia resembles totalitarian regimes or Nazi Germany;
- that technology should never be used in intelligence assessment;
- that all automated systems are inherently problematic; or
- that protecting Jewish Australians conflicts with protecting other communities.

What this submission does argue is that institutional design matters—that well-intentioned responses to genuine threats can nonetheless create structural vulnerabilities when safeguards are absent or inadequate, and that recent Australian experience demonstrates this is not a theoretical concern but a documented pattern.

The Commission's Opportunity

The Royal Commission into Antisemitism and Social Cohesion is uniquely positioned to address these concerns **before** Schedule 2 is widely implemented.

This differs fundamentally from Robodebt.

The Robodebt Royal Commission documented harm after 443,000 people had been affected, after documented suicides, after years of unlawful operation. It was a retrospective accounting of systemic failure.

This Commission has the opportunity to act prospectively.

Schedule 2 is enacted but not yet operationalized at scale. Automated systems have not yet been deployed. Contracts under section 92E's "arrange for" authority may not yet be finalized. The machinery exists in statute, but its impacts have not yet materialized across the population.

The Commission can recommend safeguards before harm occurs rather than documenting harm after it has occurred.

This is an extraordinary opportunity. Royal Commissions rarely have the chance to prevent systemic harm rather than merely accounting for it retrospectively.

What Success Would Look Like

This submission has proposed ten specific recommendations structured in two tiers. Success would not require the Commission to accept every recommendation in its precise form. Rather, success would mean:

Immediate actions (Tier 1):

- Recognition that Schedule 2 raises serious questions requiring parliamentary reconsideration with full information;
- Advice against operationalization pending Solicitor-General constitutional review and incorporation of Robodebt-consistent safeguards;
- Referrals for investigation where evidence warrants examination of process and disclosure questions.

Structural reforms (Tier 2):

- Establishment of minimum safeguards if any automated intelligence assessment framework is retained—no immunity from invalidity, notification before adverse action, human decision-making preserved, methodology transparency, independent review, scope defined by Parliament;
- Reforms to parliamentary process ensuring future automation-related legislation receives adequate scrutiny and disclosure;
- Mechanisms ensuring Royal Commission recommendations are monitored and implemented rather than ignored; and
- International benchmarking to align Australian practices with democratic best practices.

Most importantly: A clear statement from the Commission that **effective protection of Jewish Australians requires democratic accountability, not its suspension.**

The Broader Stakes: Social Cohesion

The Commission's mandate includes social cohesion—the bonds that hold diverse communities together in a functioning democracy.

Social cohesion is undermined when:

- Legislative processes appear to lack transparency, creating perceptions that Parliament was not fully informed;
- Automated systems operate without adequate safeguards, eroding trust that government power is exercised fairly;
- Some communities feel targeted while others feel inadequately protected, creating division rather than unity;
- Lessons from recent failures (Robodebt) appear ignored, suggesting accountability mechanisms don't function; and
- Democratic restraints appear suspended in the name of protection, raising questions about whether the cure creates new vulnerabilities.

Social cohesion is strengthened when:

- All communities see that government power is exercised through law, with transparency and accountability;
- Protection is achieved through means that preserve rather than undermine democratic foundations;
- Mistakes are acknowledged, learned from, and not repeated;
- No community is scapegoated or subjected to systems that treat them as inherently suspect; and
- Liberal democratic principles—equality before law, procedural fairness, independent review—are maintained even in moments of fear.

The Commission's findings on Schedule 2 will therefore affect social cohesion far beyond the specific provisions of the legislation. They will signal whether democratic accountability matters, whether Royal Commission recommendations carry weight, and whether institutional learning from failures like Robodebt shapes future governance.

A Personal Reflection

This submission opened with a personal statement about preventing anything like the Holocaust from happening again—to any group, anywhere.

That commitment is not rhetorical. It reflects a conviction that the lesson of history's greatest atrocity is not merely to protect one group, but to protect the institutional foundations that make protection possible for all.

Liberal democracy is fragile. It depends on norms, restraints, and accountability mechanisms that can erode gradually through well-intentioned responses to genuine threats. History shows that democracies rarely collapse through dramatic coups; they weaken through incremental accumulations of power that, individually, each seem justified by crisis.

The Holocaust did not begin with death camps. It began with classification systems, legal exceptions, administrative machinery that could identify and track populations, and the normalization of treating some people as inherently threatening based on identity rather than conduct.

This is why institutional design matters. Not because contemporary Australia resembles Nazi Germany—it emphatically does not—but because the structural features that enabled past atrocities can appear in new contexts, justified by different rationales, yet creating similar vulnerabilities.

Protecting Jewish Australians from antisemitism and protecting all Australians from arbitrary government power are not competing goals. They reinforce each other. A society governed by law, with robust accountability and independent review, is the best protection any minority has. A society where power accumulates without constraint, where automation replaces judgment, where immunity shields errors from review—that society becomes dangerous for everyone, eventually including those it initially claimed to protect.

The Political Lifecycle of Security Powers: A Critical Warning

This submission has focused extensively on structural parallels to Robodebt, constitutional questions, and parliamentary process. Before concluding, one further consideration merits explicit and direct statement:

The tools created by sections 92E–92G will outlast the government that created them.

The Institutional Design Question

This submission acknowledges that current architects may genuinely intend to protect the Jewish community from antisemitism. That sincerity is not questioned.

However, **institutional analysis requires examining not only current intent but foreseeable future application**.

The Commission may encounter arguments that sections 92E–92G can be safely maintained because current political circumstances favour appropriate implementation, or because certain advocacy groups maintain sufficient political influence to ensure proper use.

This submission respectfully argues that such reasoning is the precise vulnerability these provisions create.

Why Political Confidence Cannot Replace Institutional Safeguards

Political influence is not permanent. Statutory powers are.

Australia's political landscape includes concerning far-right elements. This is documented reality: individuals displaying Nazi symbols and antisemitic messaging have demonstrated at state parliament buildings. Political parties with established records of targeting minority communities based on religion, ethnicity, and immigration status are polling at levels that could translate to increased parliamentary representation and potential coalition influence.

When political power shifts to include such elements, sections 92E–92G's capabilities will be available to them:

- Algorithmic systems that assess individuals and communities for security risk
- Ministerial discretion to expand categories without parliamentary amendment
- Immunity provisions limiting judicial challenge
- Reduced notification requirements concealing the basis of adverse determinations
- Classification frameworks preventing transparency about methodology

These features do not become inoperative when governments change. They persist.

The Historical Pattern

This is not an argument about the current government's intentions. It is an argument about **institutional design across political cycles**.

History demonstrates—repeatedly and tragically—that administrative machinery created to classify and track populations based on group characteristics does not remain limited to original purposes. The machinery operates according to its structural capabilities, not the intentions of its creators.

- Tools created to protect one group become weapons against that same group under different governments
- Powers justified by one threat become applied to different targets
- Classification systems established for one purpose prove readily adaptable to others
- Immunity provisions designed to enable one type of action shield entirely different applications

This pattern is not an unintended consequence. It is the foreseeable operation of institutional tools.

The Stakes for the Jewish Community

This submission therefore respectfully argues that the Jewish community's long-term safety—and the safety of all minorities in a pluralist democracy—depends on **not creating such tools**, regardless of current political circumstances or confidence in current influence.

The question before the Commission is not:

- Whether current drafters meant well (they may have)
- Whether current political leaders support the Jewish community (they may)
- Whether advocacy groups currently have influence (they may)

The question before the Commission is:

- What happens to these tools twenty years from now, or thirty, when very different people control them?
- Can institutional safeguards in the legislation prevent misuse across political transitions?
- Does relying on political alignment rather than structural protection represent sound governance?

This submission respectfully argues the answers are: dangerous outcomes, no, and no.

The Commission's Unique Responsibility

Unlike those who enacted sections 92E–92G, the Commission is not operating under political pressure or time constraints. The Commission has the opportunity—and the responsibility—to examine these tools with institutional hindsight and political foresight.

The Commission should not take sections 92E–92G at face value.

The Commission should not assume these provisions are appropriate instruments for implementing its recommendations simply because they exist in statute.

The Commission has a duty to critically examine whether these tools create greater long-term risks than they solve, and to recommend suspension and review before operational deployment entrenches them beyond democratic reconsideration.

That examination should inform whether the Commission recommends parliamentary review and suspension of sections 92E–92G's implementation.

Final Words to the Commission

This submission has been long, detailed, and at times critical. This reflects not dismissal of the underlying problem—antisemitism—but deep concern about the quality and sustainability of institutional responses.

The question before the Commission is ultimately simple:

Will responses to antisemitism strengthen or weaken the democratic foundations on which all communities' long-term safety depends?

Schedule 2, as currently designed, raises serious questions on that measure. Those questions are grounded in evidence, recent experience, expert warnings, and established constitutional principles.

The Commission has the opportunity to answer those questions before harm occurs.

That opportunity should not be wasted.