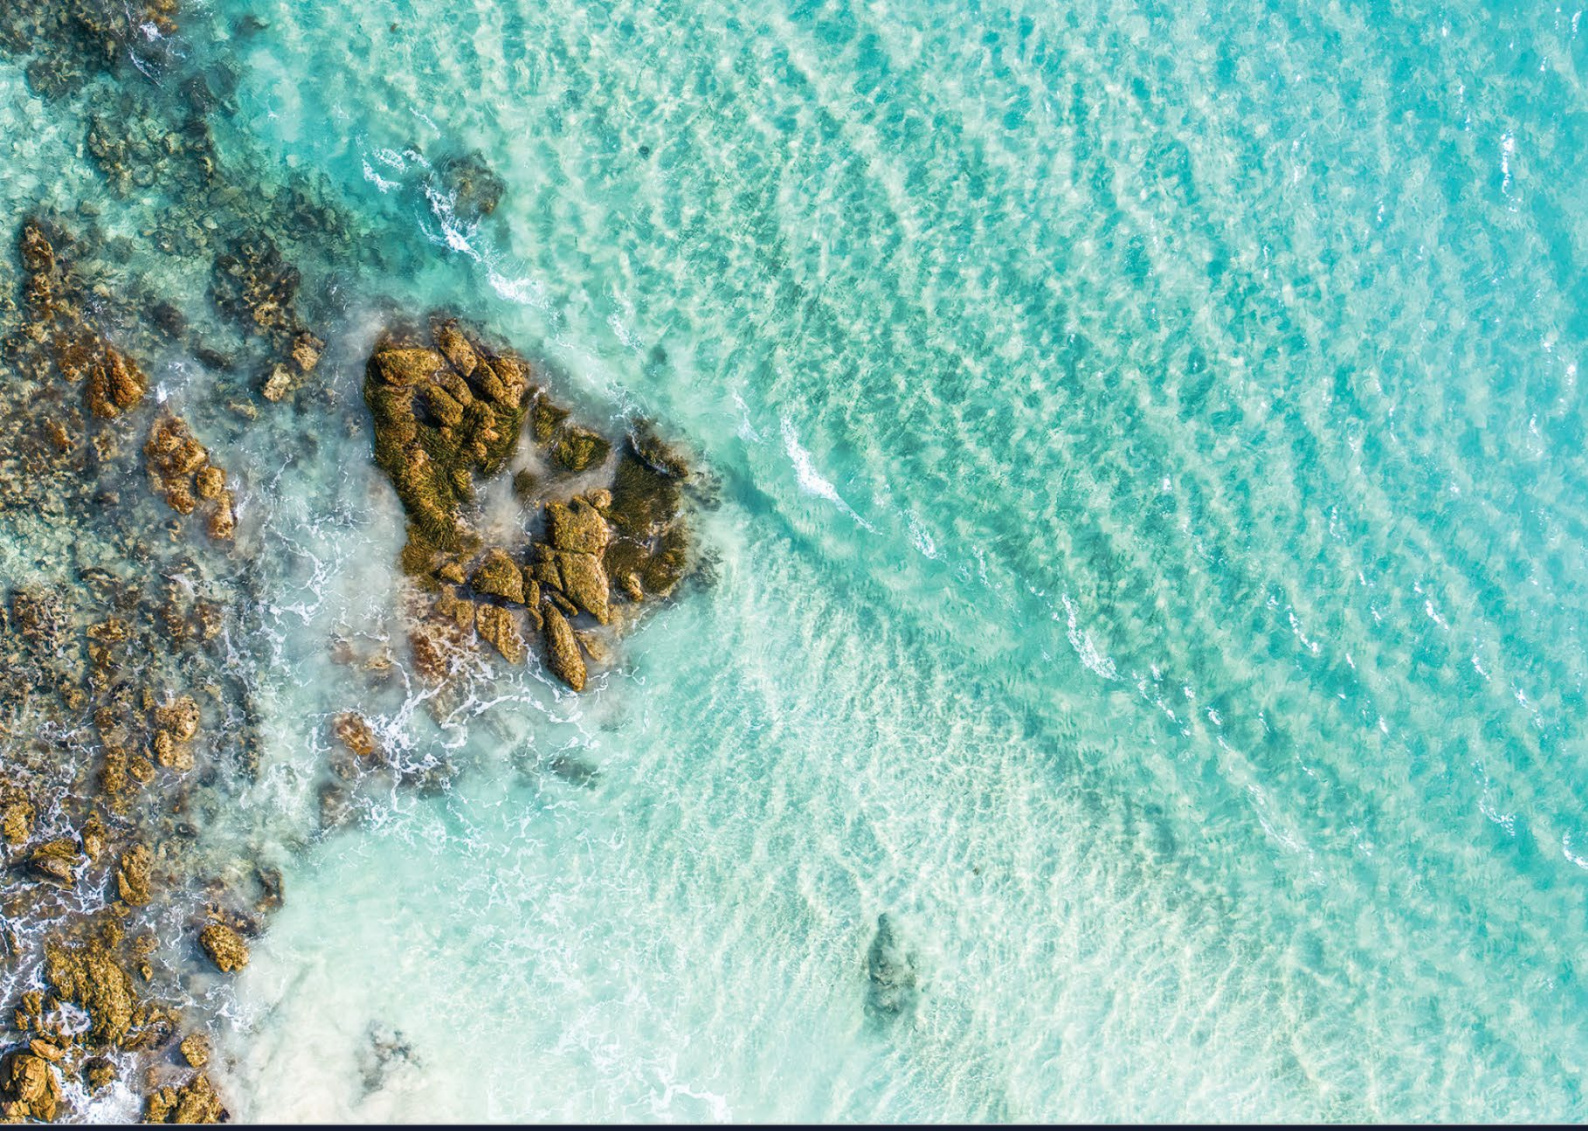

Designing Robust and Micro-Calibrated Policy

Comparative Lessons from Australian Policy Reform for Practitioners

Professor Giliberto Capano

ANZSOG Academic Fellow | University of Bologna

ANZ SOG Australia & New Zealand School Of Government

Contents

KEY INSIGHTS **3**

1. Introduction: Why Micro-Design Matters- The Framework and What It Reveals **3**

2. Case Analysis: Two Australian Policy Reforms **4**

Case One: The JRG and the Lock-In Paradox **4**

Case Two: The SMMA and the Ambition Gap **6**

3. Analysis of Insights: What the Two Cases Reveal Together **9**

4. Impact and Implementation: Five Design Principles **10**

5. Conclusion: The Cost of Not Asking **13**

6. Future Research Directions **13**

References **14**

Appendix: Technical Framework, Indices and Coding Data **15**

KEY INSIGHTS

- Micro-design choices, the technical decisions about how a policy is constructed at the instrument level, are the primary determinant of whether it achieves its objectives and whether it can be corrected when it does not.
- The Job-Ready Graduates Package (2021) looks like a simple change to university fees. It turned out to be a trap. The pricing mechanism rests on three provisions that are logically inseparable: change one and you have to change all three, and changing all three requires new legislation and new public money. The Albanese government discovered this when it tried to fix the damage: it added new funding on top instead, because undoing the original design was too expensive.
- The Social Media Minimum Age Act (2024) appears to be a landmark structural reform. It is not. Systematic analysis of its legislative architecture shows that the reform changed almost nothing in the regulatory structure it amended. Maximum political ambition produced minimum structural change, because one deliberately undefined compliance standard delegated all the operational work to actors who cannot deliver it at scale.
- Both dynamics were predictable before enactment: the JRG's lock-in and the SMMA's implementation gap are direct consequences of specific design choices, not of bad luck or poor execution. The framework presented here makes those choices visible before legislation is passed, not after.
- The empirical analysis generates five design principles: map interdependencies and their correction cost before enacting; give every open compliance standard a legally binding process for giving it operational content; inscribe behavioural outcomes in legislation with a timeframe and a measurement mechanism; match enforcement capacity to enforcement ambition; build a testing mechanism into the design for every behavioural assumption.

1. Introduction: Why Micro-Design Matters- The Framework and What It Reveals

Policies underperform for many reasons. Bad politics, insufficient resources, external shocks, poor implementation. But two of Australia's most prominent recent reforms have produced unintended consequences and structural rigidities that no amount of political will can easily undo, not because of bad politics or poor execution, but because of specific choices made in how the instruments were constructed at the level of legislative text.

This paper presents findings from an ANZSOG Academic Fellowship built around a framework designed to make those weaknesses visible before a reform is enacted, not after. The framework works by disaggregating a policy into its constituent instruments and examining each instrument along two sets of dimensions. [1]

[1] The analysis draws on two sources. First, an original instrument-by-instrument coding of the two reforms along the ten micro-design dimensions and the robustness dimensions set out in the Appendix. Second, a set of semi-structured

The first set concerns goal specifications. For each instrument, is the target population explicitly identified? Is the expected outcome stated? Is there a timeframe for achieving it? These questions determine whether a policy can be held accountable for its results. A policy that cannot answer all three has not actually committed to its own declared objective. It has announced an intention.

The second set concerns the calibration of the operational machinery. For each instrument: how strictly does it bind its targets? Does it execute automatically, or does it depend on actors choosing to act? How much discretion is left to implementers? Which organisations are responsible, and do they have the capacity to deliver? Are there monitoring mechanisms that connect performance evidence to consequences?

Together, these two sets of dimensions reveal the micro-design architecture of a reform: the concrete configuration of choices that determine whether the policy can do what it says it will do, detect when it is not doing it, and correct course before the consequences become irreversible. This architecture is almost never examined systematically in policy analysis. Political commentary focuses on macro intent. Academic evaluation focuses on outcomes after the fact. The framework focuses on the design choices that determine outcomes before they materialise. The technical details of the framework and the full coding data are presented in the Appendix.

2. Case Analysis: Two Australian Policy Reforms

Case One: The JRG and the Lock-In Paradox

What it looked like

The Job-Ready Graduates Package restructured Commonwealth higher education funding by replacing twelve discipline clusters with four and creating sharp price differentials between them. The highest student contribution rose from approximately $11,200 to $14,500 per year; the lowest fell from $6,700 to $3,900. This looked incremental: a change to numbers in a table, the kind of adjustment a future government could revisit if it did not work as planned.

Expert interviews conducted tell a more complex story. The actual design drivers were fiscal: to unfreeze Commonwealth university contributions frozen since 2017, to manage the anticipated enrolment surge from the Costello baby boom, and to do all of this with zero net increase in government spending. The declared rationale, steering students towards national-priority fields through price signals, was secondary. The fiscal constraint was primary, shaping every micro-design choice that followed.

What the analysis reveals

The analysis of 190 provisions reveals a strikingly concentrated design. The reform's entire operative weight rests on three provisions: section 30-15 (the cluster definitions), section 33-10 (Commonwealth contribution

interviews with officials and stakeholders involved in the two processes, conducted during the Fellowship. Unless attributed to a published source, the figures reported here (including staffing and calibration data) are primary findings derived from this coding and these interviews.

amounts) and section 93-10 (maximum student contribution amounts). Everything else, 62 per cent of the corpus, is definitional scaffolding.

These three provisions are logically indivisible**.** The behavioural effect depends on relative differences between clusters, not on absolute values. Modify any one without modifying the others, and the mechanism becomes operationally incoherent. This is the lock-in paradox: what appeared to be a reversible parametric adjustment was structurally irreversible from the moment of enactment.

The lock-in is confirmed by what the Albanese government did: rather than correcting the mechanism, it added new Commonwealth funding on top through the Universities Accord process. IRU (2026) modelling puts the cost of full reversal at $1.917 billion per annum.

The robustness assessment in Table 1 translates this structural diagnosis into systematic terms. Across the five dimensions that capture a design's capacity to adapt, learn and correct, the JRG scores Very Low on all but one. The exception is automaticity: the mechanism executes without anyone needing to act. From the perspective of the government that enacted it, this was a deliberate asset, a way of locking in a policy architecture that future governments could not easily unwind. It worked. But automaticity without modularity or dynamism is a one-way door: the design enforces itself regardless of what the evidence shows, and correcting it requires dismantling it entirely.

Table 1. JRG robustness assessment

Dimension	Score	What it means in practice
Modularity	Very Low	Cannot change one cluster without changing all; the three core provisions are a single indivisible block
Dynamism	Very Low	No feedback loop; no mechanism to detect or respond to underperformance
Prototyping	Very Low	No pilot; full deployment with no staged learning or correction pathway
Automaticity	High (trap)	Self-executing with no discretion: robust as a mechanism, unreachable as a problem
Polit.-Inst. Capacity	High then zero	Maximum capacity at passage; zero capacity for any subsequent government to correct it
OVERALL	**Very Low**	Durable but unadaptable: the worst combination for long-run performance

What happened

The JRG was designed to redirect student enrolments towards national-priority fields by making humanities, law and commerce more expensive and STEM, nursing and teaching cheaper. The evidence shows it did not work. IRU (2026) analysis shows that by 2024, universities received $813 million less in base funding annually; government contributed $1.18 billion less; low-SES commencements fell 9.8 per cent against a 2.2 per cent decline for non-low-SES students; in humanities, law and commerce, low-SES commencements fell 19.7 per cent. Total bachelor commencements declined 3.5 per cent. The Productivity Commission (2023) found no evidence of meaningful behavioural redirection. Daly and Lewis (2020) had predicted exactly this before enactment.

The design contained no mechanism to verify whether the price signal was redirecting student choices as intended, and no trigger for revision if it was not. Four years of evidence showing declining enrolments, widening equity gaps, and zero behavioural redirection did not prompt a correction to the JRG. They generated the Australian Universities Accord, a major independent review commissioned by the Albanese government in 2022 that called for urgent remediation of the package but stopped short of dismantling it, precisely because the cost of full reversal was estimated at $1.917 billion per annum. The reform persists not because it achieved its declared objective but because the design made undoing it more expensive than living with it.

Case Two: The SMMA and the Ambition Gap

What it looks like

The Online Safety Amendment (Social Media Minimum Age) Act 2024 requires social media platforms to take reasonable steps to prevent users under 16 from holding accounts, on pain of penalties of up to AUD $49.5 million. It was passed in eight parliamentary sitting days (introduced on 21 November 2024 and passed on 29 November 2024) and was the first legislation of its kind anywhere in the world. Implementation became mandatory on 10 December 2025, with three regulators assigned implementation roles: the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts (policy), the eSafety Commissioner (enforcement) and the Office of the Australian Information Commissioner (privacy compliance).

As Flew, Koskie and Stepnik (2025) document, the reform moved from idea to law in six months, driven by electoral competition and political bidding between parties. Speed of enactment was the priority. What the empirical analysis reveals is the cost of that priority: a design that signals maximum ambition while changing almost nothing in the regulatory architecture it amended.

What the analysis reveals

The SMMA amended an Online Safety Act that already worked. Before the reform, the Act contained well-calibrated instruments for cyberbullying, industry codes and platform standards: clear targets, designated implementing organisations, functioning monitoring mechanisms. Platforms were already required to address cyberbullying through industry codes, respond to complaints about

harmful content, and comply with Basic Online Safety Expectations that eSafety could enforce. The SMMA did not strengthen this architecture. It added a large number of definitional, procedural, immunity and transitional provisions around a single operative core, diluting the overall calibration quality of what was already there.

That operative core is section 63D (Part 4A): the one provision that carries the entire behavioural obligation of the reform. Its formal design is maximally ambitious: penalties of up to $49.5 million, a mandatory obligation, an explicit target population. Its operative design is maximally open: platforms must take reasonable steps, a standard that is deliberately undefined in primary legislation and left to non-binding guidance issued by eSafety in September 2025.

This was a conscious choice. The official documentation accompanying the legislation confirms that technology evolves too rapidly for fixed compliance definitions. The logic is defensible. The consequence is not. An undefined compliance standard only delivers results when the regulators responsible for filling it have the resources and authority to do so consistently and at scale.

The eSafety Commissioner has over forty staff on SMMA implementation. The OAIC, responsible for enforcing the privacy compliance requirement under section 63F, which requires that all age-verification data be used only for verification and then destroyed, has one to two staff. The OAIC spent seven years of investigations and litigation against Meta over the Cambridge Analytica matter before settling in December 2024 for $50 million, with no admission of liability from Meta on whether Australian privacy law had been breached. One regulatory official confirmed in interview: privacy enforcement against big tech is nearly impossible with existing powers and resources. The design requires both regulators to function. It adequately resources only one.

Table 2 presents the full robustness assessment. Unlike the JRG, the SMMA has genuine adaptive capacity: its rulemaking powers and two-year statutory review give it structural flexibility that the JRG entirely lacks, which is why it scores Medium on Dynamism. But on the dimensions that determine whether the core obligation can actually be enforced, the picture is consistently weak. Automaticity is Very Low because nothing happens without the Commissioner choosing to act: enforcement is entirely discretionary, not structural. Monitoring is Weak because compliance data is held by the platforms themselves, not by the regulators, which means the regulator depends on self-reporting to know whether platforms are complying. Redundancy is Low because there is no alternative enforcement pathway if the primary route fails or is legally contested. The SMMA's adaptive mechanisms are real and represent a genuine design strength. But adaptability is only valuable if the actors responsible for implementation can exercise it. On that question, the asymmetry documented above, between an eSafety Commissioner with over forty dedicated staff and an OAIC with one to two, is not an administrative detail. It is the central operational constraint of the entire design.

Table 2. SMMA robustness assessment

Dimension	Score	What it means in practice
Automaticity	Very Low	Nothing happens without the Commissioner choosing to act; enforcement is entirely discretionary
Monitoring	Weak	Compliance data held by platforms, not regulators; self-reporting is the primary mechanism
Polycentricity	Partial	Three regulators with a formal MOU but asymmetric capacity: eSafety 40+ staff, OAIC 1 to 2
Redundancy	Low	No alternative enforcement pathway if the primary route fails or is contested
Dynamism	Medium	Rulemaking powers and two-year statutory review create structural adaptability
OVERALL	**Low / Partial**	More adaptable than JRG but dependent on actors who are not equally resourced

Early implementation

By mid-December 2025, platforms had restricted access to 4.7 million accounts identified as belonging to under-16s (eSafety Commissioner, 2026a). This is an early indication of compliance, not evidence of effectiveness: as of February 2026, under-16s remained able to access some platforms, and young people for whom online connection is protective rather than harmful were experiencing social isolation as an unintended consequence of the restrictions (Bravehearts, 2025; Robinson et al., 2025). By March 2026, eSafety had committed to enforcement decisions by mid-2026 (eSafety Commissioner, 2026b). What counts as reasonable steps is being determined case by case in enforcement proceedings. The two-year statutory review will assess whether the compliance process is functioning, not whether under-16s have actually left social media.

3. Analysis of Insights: What the Two Cases Reveal Together

The two reforms were designed in opposite ways and have produced opposite structural problems.

The JRG fixed its core mechanism in primary legislation with no room for adjustment. The three fee provisions at the heart of the reform are mutually dependent: the price signal that was supposed to redirect student choices depends on the relative differences between clusters, not on absolute values. Change one provision and the others become incoherent. Change all three and you need new legislation and new public money. The design enforces itself automatically, with no discretion left to anyone, which means it also corrects itself never.

The SMMA did the opposite. It created a powerful obligation on paper and left its content entirely open. Platforms must take reasonable steps. What those steps are is determined case by case through regulatory enforcement and non-binding guidance. The design depends entirely on regulators to give the obligation operational meaning, and those regulators do not have equal capacity to do so.

Table 3 sets out both cases across the dimensions that matter most for practitioners.

Table 3. Two reforms, two pathologies

Design dimension	JRG	SMMA
Political presentation	A targeted reform to redirect student enrolments towards national-priority fields through price signals	A world-first reform to protect children from online harm by barring under-16s from social media
What the analysis shows	A pricing mechanism that cannot be corrected without dismantling it entirely	A regulatory obligation whose compliance content remains undefined more than a year after enactment
How the core obligation works	Three mutually dependent fee provisions: changing one requires changing all, at a cost of $1.917 billion per annum in full reversal	One provision requiring platforms to take reasonable steps: a standard determined case by case through regulatory enforcement and non-binding guidance
Timeframe for the behavioural objective	None in primary legislation	None in primary legislation: the statutory review tracks whether the compliance process is

		functioning, not whether under-16s have left social media
Can it be corrected?	No: correction requires new primary legislation and new public spending	In principle yes: rulemaking powers and the statutory review allow adjustment. In practice: the capacity gap between eSafety and the OAIC is the binding constraint
Main structural risk	The mechanism enforces itself regardless of whether it achieves its declared objective	The mechanism only works if regulators act, and they do not have equal capacity to do so

Despite their differences, both reforms share one critical design flaw: neither defines or specifies what success looks like. The JRG contains no provision stating when students should have redirected their field choices, by how much, or what should happen if the price signal proves ineffective. The SMMA contains no provision stating when under-16s should have left social media, what level of platform compliance would constitute achievement of the reform's goal, or what should trigger a redesign if enforcement at scale proves impossible.

This means that both reforms can be fully in force, formally reviewed and declared compliant while their declared goals remain unachieved, and nothing in either design requires anyone to respond. Governments that do not write their objectives into law as measurable, time-bound obligations retain the freedom to claim that a reform is working regardless of the evidence. Both the JRG and the SMMA were designed that way.

4. Impact and Implementation: Five Design Principles

These principles are not abstract recommendations. Each one identifies a specific question that was not asked during the design of the JRG or the SMMA. Had these questions been asked, the reforms would have been designed differently. Practitioners can apply them to any reform before it reaches Parliament. This paper provides the research foundation for a forthcoming diagnostic framework that will structure these principles to a method for assessing policy robustness.

1. **Before finalising the design, ask what it would cost a future government to change it.**
   The team that designed the JRG did not ask this question. The three provisions at the heart of the reform are mutually dependent: the provision that defines the four funding clusters, the provision that sets the government contribution for each cluster, and the provision that caps what students can be

charged per cluster. The price signal that was supposed to redirect student choices operates through differences across clusters, not through any single provision. Changing one without changing the others makes the mechanism incoherent. Changing all three requires new legislation and new public money. The current cost of full reversal is estimated at $1.917 billion per annum. No one calculated this at the design stage.

Before any reform is finalised, map the dependencies between its core provisions. If changing one requires changing all, make that a conscious political decision with a documented rationale. Discovering the constraint only after years of evidence that the reform is not working, as the JRG demonstrates, is not an acceptable design outcome.

2. **If you cannot define compliance in the law, specify who will define it, when, and with what authority.**
   The SMMA requires platforms to take reasonable steps to prevent under-16s from holding accounts. This standard is deliberately open: the government judged that defining compliance precisely in primary legislation would make the law obsolete as technology evolved. That judgment may be correct. Its consequence is that the content of the central obligation of the reform is not in the law. It is being determined case by case through regulatory enforcement, guided by non-binding guidance that eSafety issued in September 2025, ten months after the legislation was passed.

   Every reform that leaves a compliance standard open needs to specify, in the legislation itself or in binding rules made under it: who has the authority to give that standard operational content, through what process, by when, and what happens when regulated parties disagree with the regulator's interpretation. The SMMA answers none of these questions. Non-binding guidance that can be contested in enforcement proceedings is not a compliance framework. The strength of the reform's central obligation is currently determined by how those proceedings develop, not by what the law says.

3. **Write into the law what you are trying to achieve, by when, and how you will know if it is working.**
   Both reforms declared behavioural objectives in political statements that appear nowhere in primary legislation as obligations with a timeframe and a way of measuring whether they are being met. The JRG was designed to redirect student enrolments towards national-priority fields. No provision states when that redirection should have occurred, what level of change would constitute success, or what should happen if the price signal proves ineffective. The SMMA was designed to protect children from online harm by removing them from social media. No provision states when under-16s should have left the platforms, what proportion of platform compliance constitutes achievement of the reform's goal, or what should trigger a fundamental rethink if enforcement proves insufficient.

   A review mechanism that measures how many notices a regulator has issued is not outcome accountability. It is process accountability. Before any reform is enacted, the responsible policy team should be able to answer four questions: what specific change in behaviour are we trying to produce? By when? How will we measure it? What happens if the measurement shows we are not on track? A government that cannot answer these questions before Parliament votes has not committed to achieving what it says the reform is for.

4. **Before enacting the ambition, verify that the institutions responsible for delivering it have the capacity to do so.**
   The SMMA requires three institutions to give it operational effect. The eSafety Commissioner enforces the obligation on platforms to prevent under-16 accounts. The Office of the Australian Information Commissioner enforces the privacy compliance requirement, which mandates that all age-verification data be used only for that purpose and then destroyed. The Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts holds overall policy responsibility. eSafety has over forty staff dedicated to SMMA implementation. The OAIC has one to two. The OAIC spent seven years on investigations and litigation against Meta over the Cambridge Analytica matter before settling in December 2024 for $50 million, with no admission of liability from Meta on whether Australian privacy law had been breached. Privacy enforcement against global technology companies is structurally difficult with existing powers and resources.

   This was not a problem that emerged after enactment. It was a predictable consequence of assigning responsibilities without assessing whether the responsible institutions could carry them out. Before any regulatory reform is finalised, an explicit assessment should be produced for each implementing institution: what does this legislation require of them, do they currently have that capacity, and if not, what is the concrete plan for building it before the legislation takes effect? That assessment should inform the design, not follow it.

5. **For every assumption about how people or organisations will behave, build in a way to find out if you are wrong.**
   Both reforms rested on assumptions about behaviour that were contested before enactment. The JRG assumed that making humanities and law more expensive and STEM and nursing cheaper would redirect student choices towards national-priority fields. Daly and Lewis (2020) had demonstrated, before the reform passed, that evidence from previous changes to student fees suggested the effect on field choice would be small. The SMMA assumes that restricting access to social media platforms will improve the wellbeing of children under 16. The research evidence on this question remains genuinely uncertain: the associations between social media use and adolescent harm are real in some contexts, but the causal evidence for the benefits of age restrictions specifically is limited (Livingstone and Sylwander, 2026; Thimm-Kaiser and Keyes, 2025; Robinson et al., 2025).

Neither reform contains anything that would tell the responsible institutions whether the behavioural assumption is holding, and neither contains a process that would be triggered if the evidence showed it was not. This is not about requiring certainty before legislating. Governments must act under uncertainty. It is about acknowledging uncertainty and building in the capacity to respond to it. Before any reform that depends on a behavioural assumption is finalised, three things should be in the design: a measurable indicator that will show whether the assumption is holding in practice; a threshold at which that indicator triggers a formal review; and a body with clear authority to propose design changes when the review finds the assumption is not being borne out. Neither the JRG nor the SMMA has any of the three.

5. Conclusion: The Cost of Not Asking

The JRG and the SMMA have produced different problems at different speeds. The JRG's consequences are now documented and irreversible within the current design. The SMMA's are still emerging. What the analysis reveals in both cases is the same: the design choices that determined these outcomes were made at the drafting stage, they were identifiable before enactment, and they were not examined systematically.

The five questions behind the five principles in this paper are strategic questions for senior policy leaders, not technical questions for legal drafters. The evidence from both cases suggests they are not being asked routinely. That is what this paper is about.

The stakes extend well beyond Australia. The SMMA model is being actively considered in Ireland, New Zealand, Singapore and across the European Union (OECD, 2025). If the undefined compliance standard, the asymmetric regulatory architecture and the absence of outcome accountability travel with the legislative model, every jurisdiction that adopts it inherits the same structural weaknesses.

The JRG demonstrates something equally important: that a reform can persist for years, survive changes of government and accumulate documented evidence that it is not working, and remain structurally impossible to correct, because no one at the design stage asked what correction would require.

Applying the analytical framework developed in this research to the systematic coding of both reforms makes these design choices visible before legislation is passed, not after. What this paper cannot produce is the decision to use that framework before a bill goes to Parliament. That decision belongs to the people who read it.

6. Future Research Directions

The next step is to refine and publish the companion diagnostic framework for use during policy development, legislative design and implementation planning. Future work could also test the micro-design framework across additional policy domains, examine how the diagnostic questions can be embedded earlier in Cabinet and legislative development processes, and explore how policy teams can use the framework before reforms are enacted.

References

Ansell, C., Sorensen, E., Torfing, J., & Trondal, J. (2024). Robust governance in turbulent times. Cambridge University Press.

Bravehearts. (2025). Social media age restriction and young people [Briefing paper]. https://bravehearts.org.au

Capano, G., & Howlett, M. (2024). The devil in the details: Measuring micro-design change and policy innovation. Policy Sciences. Advance online publication. https://doi.org/10.1007/s11077-024-09548-x

Capano, G., & Woo, J. J. (2017). Resilience and robustness in policy design: A critical appraisal. Policy Sciences, 50(3), 399-426. https://doi.org/10.1007/s11077-016-9273-x

Daly, A., & Lewis, P. (2020). The proposed Job-Ready Graduate Package: A misguided arrow missing its target. Australian Journal of Labour Economics, 23(2), 231-250.

eSafety Commissioner. (2026a, January 16). Platforms restrict access to 4.7 million under-16 accounts across Australia [Media release]. https://www.esafety.gov.au/newsroom/media-releases/platforms-restrict-access-to-47-million-under-16-accounts-across-australia

eSafety Commissioner. (2026b, March 31). Social media minimum age compliance update. https://www.esafety.gov.au/sites/default/files/2026-03/SocialMediaMinimumAgeComplianceUpdateMarch2026.pdf

Flew, T., Koskie, T., & Stepnik, A. (2025). Digital policy as problem space: Policy formation, public opinion, and Australia's Online Safety Amendment (Social Media Minimum Age) Act 2024. Media International Australia. https://doi.org/10.1177/1329878X251406315

Innovative Research Universities. (2026, March). Impacts of the Job-Ready Graduates policy and options for reform. IRU.

Livingstone, S., & Sylwander, K. R. (2026). Conceptualizing age-appropriate social media to support children's digital futures. British Journal of Developmental Psychology, 44, 300-312. https://doi.org/10.1111/bjdp.70006

Norton, A. (2023, November). Job-Ready Graduates 2.0: The Universities Accord and centralised control of universities and courses (Analysis Paper 58). Grattan Institute.

OECD. (2025, June 12). Too young to scroll? Why governments are cracking down on social media age limits. OECD Blog. https://www.oecd.org/en/blogs/2025/06/too-young-to-scroll-why-governments-are-cracking-down-on-social-media-age-limits.html

O'Donnell, L. (2025). Australia's social media minimum age laws. Alternative Law Journal, 50(4), 305-307.

Productivity Commission. (2023). *Advancing prosperity: 5-year productivity inquiry* (Report no. 100, Vol. 8: From learning to growth). Australian Government.

Robinson, J., La Sala, L., & Harrison, V. (2025). Australia's social media age limit: A "seatbelt moment" or a missed opportunity for a nuanced approach? Crisis, 46(5), 245-253. https://doi.org/10.1027/0227-5910/a001014

Thimm-Kaiser, M., & Keyes, K. M. (2025). US state policies regarding social media: Do policies match the evidence? The Milbank Quarterly, 103(S1), 337-365. https://doi.org/10.1111/1468-0009.70021

Yong, M., Coelli, M., & Kabatek, J. (2023). University fees, subsidies and field of study (Melbourne Institute Working Paper No. 11/23). Melbourne Institute.

Appendix: Technical Framework, Indices and Coding Data

A1. The micro-design framework

The framework assesses each instrument along ten dimensions in two groups.

Table A1. Target specifications and calibrations: coding rules

	Dimension	Coding rule
G1	Target population	Binary (0/1): 1 if the provision explicitly designates the population it targets
G2	Expected outcome	Binary (0/1): 1 if the provision states an expected outcome of the intervention
G3	Timeframe	Binary (0/1): 1 if the provision specifies a timeframe for achieving the aim
C1	Primary recipients	Binary (0/1): 1 if the provision explicitly identifies primary recipients
C2	Intrusiveness	Ordinal 0 to 3 (0 absent, 1 low, 2 medium, 3 high). Scaled: (mean/3) x 100
C3	Additional resources	Binary (0/1): 1 if the provision allocates additional resources for implementation
C4	Organisation in charge	Binary (0/1): 1 if the provision designates the implementing organisation
C5	Discretion (inverted)	Ordinal 0 to 3. Inverted and scaled: ((3-mean)/3) x 100. Lower discretion = better calibration
C6	Accountability rules	Binary (0/1): 1 if the provision includes explicit accountability mechanisms
C7	Monitoring	Ordinal 0 to 2 (0 absent, 1 moderate, 2 strong). Scaled: (mean/2) x 100

A2. Robustness dimensions

Table A2. Robustness dimensions and diagnostic questions

Dimension	Diagnostic question
Modularity	Can individual components be changed without requiring changes to all others?
Dynamism	Are feedback mechanisms built in that allow the design to detect and adjust to performance problems?

Prototyping	Was iterative learning incorporated through pilots or staged implementation?
Scalability	Can the design respond to unexpected changes in demand or context?
Polycentricity	Does multi-actor coordination function across the implementation network?
Redundancy	Are there backup mechanisms if the primary implementation route fails?
Bounded Autonomy	Does the design allow appropriate discretion at implementation level without losing coherence?
Automaticity	Does compliance occur without continuous actor intervention?
Polit.-Inst. Capacity	Do the responsible institutions have the authority, resources and expertise to steer and enforce?
Strategic Polyvalence	Can the design shift direction if the context or the evidence base changes significantly?

ANZSOG

anzsog.edu.au