Regulatory sandboxes fail because they get the default rule wrong

Submission to the Review of the Enhanced Regulatory Sandbox

Authors: Associate Professor Darcy W.E. Allen
Professor Chris Berg

Contact: Darcy Allen, darcywilliamallen@gmail.com

Date: 21 January 2026

Dear Maha El Dimachki,

We submit this response to Australia's Independent Review of the Enhanced Regulatory Sandbox (ERS) as academics specialising in innovation policy, regulatory design, and law and economics.[1] Our work examines how innovation interacts with, and shapes, institutions.[2] We have also published several books and other research on regulatory reform tools in Australia.[3] Our submission draws on this expertise, including in the attached working paper on regulatory sandboxes.[4]

We argue that the ERS underperforms not because of eligibility thresholds or administrative frictions, but because it embeds the wrong "default rule". Like most sandboxes, the ERS is a closed, permission-based sandbox that treats experimentation as unlawful unless approved. This design choice predictably generates low uptake, adverse selection, and weak learning. **Effective reform requires reversing the baseline default rule presumption — from prohibition to permission — through an open safe-harbour design.**

Our submission focuses on Consultation Paper Question 1 ("What are the benefits and limitations of the current ERS?") and Question 6 ("How can the current eligibility criteria be improved to increase participant uptake?"), while indirectly contributing to other questions. Our attached working paper,

[1] We are academics at RMIT University but make this submission in our personal capacities.

[2] For instance, see Allen, D. W. E., Berg, C., and Potts, J. (2025). *Institutional Acceleration: The Consequences of Technological Change in a Digital Economy*. Cambridge University Press; Allen, D. W. E., Berg, C., and Lane, A. M. (2019). *Cryptodemocracy: How blockchain can radically expand democratic choice*. Lexington.; Allen, D. W. E., Lane, A. M., & Poblet, M. (2019). "The governance of blockchain dispute resolution". *Harvard Negotiation Law Review*, 25, 75; Lane, A.M., Allen, D. W. E. and Berg, C. (2024) "Towards Legal Recognition of Decentralised Autonomous Organisations". *Australian Business Law Review* 52(2): 96-116.

[3] See Allen, D.W.E. and Berg, C. (eds.) (2018). *Australia's Red Tape Crisis: The Causes and Costs of Over-Regulation*, Connor Court, Brisbane, Australia; Allen, D.W.E., Berg, C., Lane, A.M., and McLaughlin, P.A. (2021). "The Political Economy of Australian Regulatory Reform" *Australian Journal of Public Administration* 80(1): 114-137.

[4] Allen, D.W.E., Berg, C., and Lane, A.M. (2026). "From Closed to Open Regulatory Sandboxes." Working Paper available at SSRN.

"From Closed to Open Regulatory Sandboxes" provides the theoretical framework and evidence for our responses to these questions, while in this submission we summarise the key arguments.

Default and altering rules

Like most sandboxes, the Enhanced Regulatory Sandbox (ERS) is closed and permission-based. Experimentation is unlawful unless and until a regulator grants permission. Entry is mediated through application, discretionary assessment, and time-limited approval. This design choice is not incidental. It reflects a specific allocation of legal default rules: prohibition by default, with innovators required to contract around that prohibition through an administrative process. This default-rule choice largely explains the ERS's persistent low uptake. But low participation is not a puzzle to be solved through marginal process improvements. It is the predictable equilibrium outcome of a sandbox that embeds a prohibitory baseline and relies on costly altering rules to escape it.

Under the ERS, firms wishing to experiment must opt out of the default by navigating an application process that includes documentation, legal analysis, uncertain timelines, and discretionary judgment. From a law and economics perspective, the application process functions as an "altering rule". That is, the sandbox application process is the mechanism through which parties contract around an unfavourable default. Where altering rules are costly, uncertain, or discretionary, defaults become "sticky" and persist even when they are inefficient.

The consequences of this design are observable. Closed sandboxes generate adverse selection where firms that are already close to compliance, well-resourced, and legally sophisticated are most willing to bear the costs of application and uncertainty. Firms operating at the frontier — those facing genuine regulatory ambiguity and for whom sandbox relief would be most valuable — are rationally screened out. The sandbox therefore attracts those least in need of experimental space, while excluding those for whom it was ostensibly designed.

This result is not a failure of execution. Shifting the altering rules — such as streamlining forms, accelerating approvals, or clarifying guidance — may reduce frictions at the margin, but they do not fundamentally change how legal uncertainty is allocated. As long as experimentation remains unlawful by the default rules, innovators must continue to pay the full transaction costs of opting out ex ante. The ERS will remain closed in substance despite becoming more efficient in procedure.

If the objective of the Review is to increase participation by firms operating under genuine uncertainty — particularly in a technological environment characterised by novelty and rapid change — then the problem to be addressed is not eligibility thresholds or administrative complexity, but the default rule itself. Without changing the baseline presumption from prohibition to permission, the ERS will continue to exhibit low uptake regardless of how efficiently the permission process is administered.

Limitations of sandboxes

The central limitation of the ERS is that it treats experimentation as an exception requiring permission, rather than as a default activity conducted under regulatory oversight. This design choice produces a set of predictable limitations that compound and reinforce one another.

First, the ERS front-loads regulatory uncertainty. Firms must translate incomplete, evolving knowledge about novel products into regulator-legible applications before product-market fit, technical feasibility, or risk profiles are known. This requirement favours firms with established legal and compliance capacity and disadvantages genuinely experimental activity. The cost is not simply administrative. It is the cost of being required to specify answers to questions that cannot yet be answered with confidence.

Second, the ERS relies on discretionary selection under political and reputational constraints. Regulators bear the consequences of sandbox failures, while the benefits of experimentation are diffuse and long-term. This incentive structure biases selection toward low-variance, compliance-adjacent business models. As a result, sandbox participation skews toward firms already on conventional regulatory trajectories, rather than toward frontier activity where learning value is highest.

Third, the ERS exposes firms to strategic disclosure and hold-up risk. Participation requires early revelation of business models, technical architectures, and legal interpretations. This disclosure occurs before market validation and can trigger bespoke regulatory constraints or supervisory expectations that persist beyond the sandbox period. Anticipating this risk, many firms rationally avoid participation altogether.

Finally, time-limited relief without credible scaling pathways suppresses participation. Firms must anticipate meeting full licensing requirements at the end of the sandbox period, often without clarity as to how or when underlying regulatory frameworks will adapt. This "exit cliff" further concentrates participation among firms already close to compliance and discourages experimentation by firms exploring genuinely new organisational or technological forms.

These limitations are the mechanisms through which a closed sandbox remains closed in practice. Taken together, they explain why low participation is a stable outcome of the ERS rather than a temporary implementation problem. Addressing them requires more than marginal procedural reform. It requires reconsidering the baseline assumption that experimentation must be individually authorised in advance.

Combinatorial innovation compounds the problem

These structural problems compound as innovation becomes increasingly combinatorial. As we outline in our recent book, *Institutional Acceleration: The Consequences of Technological Change in a Digital Economy,* this combinatorial nature of innovation leads to institutional acceleration as new technologies combine and create competing rule systems.[5] And as we explain in our working paper:

> innovation is becoming more open, global, and combinatorial… Frontier technologies increasingly arrive as bundles of capability that combine multiple layers of technology and governance, and that cut across sectoral and legal boundaries.[6]

[5] Allen, D. W. E., Berg, C., & Potts, J. (2025). *Institutional Acceleration: The Consequences of Technological Change in a Digital Economy*. Cambridge University Press.
[6] Allen et al. 2026, p. 10.

For instance, an AI agent executing crypto-denominated transactions may implicate financial services licensing, custody rules, market conduct standards, corporate law, and AML controls. Narrow relief from a single licensing requirement provides negligible value when exposure under other regimes remains sufficient to deter activity. A sectoral sandbox design cannot accommodate distributed legal exposure. Partial relief is barely superior to no relief when the probability that at least one applicable regime prohibits activity approaches certainty.

Attempts to broaden sandbox scope elsewhere — such as Utah's all-industry model and Senator Cruz's SANDBOX Act — are steps in the right direction, but without altering the underlying default rule they do not solve this problem. Rather, they make closed sandboxes marginally faster or broader while still preserving the prohibitory default and its associated transaction costs.

Sandbox participation as an incomplete contracting problem

Every regulatory system must answer a baseline question: *what is the legal status of novel activity that does not fit existing categories*? Current sandboxes overwhelmingly adopt a closed approach because activity is prohibited by default and permitted only through discretionary application. Firms must contract around the prohibitory baseline by negotiating individualised permission with regulators.

This creates a particular kind of economic problem. Drawing on Oliver Williamson's transaction cost economics and Oliver Hart's incomplete contracting framework, **we show this governance structure inherits the pathologies of incomplete contracts under uncertainty.** When parties cannot specify all future contingencies in advance (which is inevitable with genuinely novel technologies), they must rely on discretionary adaptation as circumstances change. In normal commercial contracts, this creates risks of fragility, renegotiation, and hold-up. But those risks are managed by contract law, courts, and the ability of parties to walk away or seek remedies when the other party acts opportunistically.

Permission-based sandboxes have none of these safeguards. As we argue in our paper:

> regulatory discretion functions as the analogue of residual control: conditions may be revised, obligations reinterpreted, or permission withdrawn as new information emerges… The arrangement therefore inherits the transaction-cost pathologies of incomplete contracting under uncertainty… without the private-law infrastructure that disciplines opportunism.

Put simply, regulators simultaneously serve as counterparty (they grant permission), interpreter (they decide what the rules mean), and enforcer (they determine whether to revoke permission). Firms have little practical remedy when regulators change course. This concentration of power makes the relationship unstable and unpredictable, exactly when innovators need stability to justify investment.

This is fundamentally misaligned with innovation under uncertainty. Innovators typically possess superior information about how novel technologies work than regulators do. Yet when products are genuinely new, the cost of specifying in advance how they should be regulated is extremely high. Under these conditions, monitoring what firms actually do and holding them accountable for harm

ex post (after the fact) is a more effective regulatory strategy than trying to predict and approve everything ex ante (before it happens).

From our perspective, low ERS participation is not a puzzle but an equilibrium outcome of closed sandbox design operating exactly as the theoretical framework predicts. Rational innovators avoid a system where the costs are high, the outcomes are uncertain, and the regulator holds all the cards.

Open sandboxes as default-in safe harbours

We propose a structural reform to create open sandboxes implemented as default-in safe harbours:

> under such a design, the regulator specifies a domain of activity, objective eligibility criteria, participation conditions (such as registration, disclosure, and reporting), limits on scale or scope, and grounds for revocation. Firms that satisfy these conditions receive protection from enforcement of specified licensing or authorisation requirements within that domain.[7]

For instance, in an open sandbox:

- Eligible firms are presumed protected from enforcement of specified authorisation requirements by default
- Protection applies subject to registration, disclosure, predefined conditions, and revocation powers
- The baseline presumption is permission, not prohibition
- Regulatory oversight operates through ex post monitoring rather than ex ante gatekeeping

This reverses the burden of proof. Rather than innovators demonstrating ex ante that they should be allowed to act, regulators must demonstrate ex post that protection should be withdrawn based on observed conduct.

There are several reasons why such a shift from closed permissioned sandboxes to open sandboxes might address some of the concerns we outlined previously. Open sandboxes could:

- Eliminate application costs as a barrier to entry. Firms meeting objective eligibility criteria can begin experimenting immediately, without discretionary approval or strategic disclosure before market validation.
- Bundle protection for combinatorial innovation by specifying protection across multiple authorisation domains simultaneously, matching the structure of innovation itself rather than forcing firms to disentangle exposure regime by regime.
- Shift regulators to evaluate actual conduct and risk profiles rather than speculating about uncertain futures, aligning oversight with information availability.

Our submission will not fit neatly with more incremental changes proposed by others (e.g. streamlined applications, accelerated approvals, expanded sectoral coverage, deemed approval mechanisms). Yet we do not believe the ERS can be made effective through incremental adjustment. If the objective is to create genuine infrastructure for experimentation under uncertainty, then the baseline presumption of legality must change.

[7] Allen et al. 2026, p. 20

Implementation challenges

The implementation of open sandboxes will be hard. We identify several challenges:

1. **How can we define eligibility without discretion?** Negative definition (protection applies where law is ambiguous) reduces but does not eliminate interpretive judgement. The critical difference is that presumption favours inclusion and burden shifts to regulators.
2. **How can we coordinate across regulators?** Where institutional authority permits, a single enabling instrument can specify protection across domains.
3. **What are the exit pathways?** Safe harbour participation may inform regulatory adaptation, creating new categories or modifying existing ones. Alternatively, safe harbours may operate as durable alternatives rather than temporary exceptions.

These challenges are not unique to default-in approaches. Closed sandboxes face the same problems, often more acutely, while enabling far less experimentation.

Concluding remarks

The persistent underutilisation of the ERS, and indeed many sandboxes globally, reflects a structural design problem rather than an implementation failure. Sandboxes are typically closed and permission-based, embedding a prohibitory default rule. Experimentation is unlawful unless approved. This design predictably generates high transaction costs, adverse selection, and low participation, particularly among firms operating under genuine regulatory uncertainty. Incremental reforms — such as streamlining applications, expanding eligibility, or accelerating approvals — cannot resolve this problem. As long as experimentation remains prohibited by default, innovators must bear the full costs of opting out ex ante. The result is a sandbox that systematically favours compliance-ready firms and screens out frontier experimentation.

The Review therefore faces a clear design choice. If the objective is to create regulatory infrastructure that supports experimentation and learning under uncertainty, the baseline presumption of illegality should change. Open sandboxes implemented as default-in safe harbours offer a viable alternative. They preserve regulatory oversight and enforcement authority while reallocating uncertainty away from innovators and toward ex post monitoring based on observed conduct.

Reframing sandboxes as legal infrastructure rather than discretionary privileges would better align regulatory design with the realities of combinatorial innovation in an open, complex digital economy. Without a change in default rule, the ERS is unlikely to achieve its stated objectives regardless of procedural refinement.

Regards,

Associate Professor Darcy W.E. Allen
Professor Chris Berg