

AJRC

Australia–Japan Research Centre

Commitment under recognition uncertainty: how fossil-market flexibility delays or accelerates green investment

AJRC Working Paper 01/05/2026
May 2026

Hiroshi Matsushima¹

The Australian National University
Australia–Japan Research Centre

Abstract

Cross-border trade in certified green commodities has stalled despite falling costs and substantial policy commitments, with green hydrogen the most visible case. I show that what holds commitment back in these markets is recognition uncertainty over whether certified output will count as “green” under the importing jurisdiction’s rules. Most existing frameworks compare supply sources with that recognition problem already resolved. I instead study when buyers make an irreversible bilateral commitment while that question remains unsettled and flexible fossil supply remains available. In that setting, incumbent flexibility acts as a “bridge” that accelerates green commitment when recognition is clear, but as a “destination” that sustains delay when recognition is uncertain. Moreover, for certified goods with near-zero embedded emissions, reducing recognition uncertainty accelerates commitment more than additional carbon-price support, all the more so where fossil-market flexibility is greater. In an Australia–Japan calibration where LNG is the incumbent market and green hydrogen is the certified supply chain, clearer recognition has shifted incumbent LNG toward a “bridge” role between 2019 and 2024. Among the instruments considered, mutual recognition brings commitment forward the most.

¹ Hiroshi Matsushima, Crawford School of Public Policy, The Australian National University, Canberra 2601, Australia. Email: hiroshi.matsushima@anu.edu.au

Keywords

Green hydrogen, eligibility uncertainty, certification, bilateral investment, real options, mutual recognition, LNG

JEL Codes

F13, F18, Q42, Q48

Suggested Citation:

Matsushima, H., 2026. AJRC Working Paper, May 2026. Australia-Japan Research Centre, Crawford School of Public Policy, The Australian National University.

Address for correspondence:

(E) ajrc@anu.edu.au

ISSN 0 728 8409

ISBN 978-0-86413-017-4

The Australia-Japan Research Centre (AJRC) conducts research to explore and improve understanding of the economies and economic policy processes in Australia and Japan and both countries' strategic interests in the Asia Pacific economy.

The AJRC Working Paper Series aims to provide a forum for the exchange of topical research and contains the latest research and analysis on the Japanese economy, as well as the political economy, regional integration and trade. The views expressed in AJRC Working Papers are those of the individual authors and do not represent the views of the Australia-Japan Research Centre, the Crawford School, or the institutions to which authors are attached.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

1 Introduction

The global imperative to limit warming to 1.5°C has made cross-border trade in certified green commodities (hydrogen, ammonia, green steel) a focal element of climate policy (IPCC, 2022; International Energy Agency, 2025a). Certification regimes and border measures now condition market access on recognition—on whether a project’s output will be treated as “green” under the importing jurisdiction’s rules. Yet despite sharp cost declines and substantial public commitments, final investment decisions (FIDs) remain rare: only about 9 percent of announced low-emissions hydrogen capacity has reached FID (International Energy Agency, 2025a).¹ Developer surveys consistently rank recognition uncertainty among the most important barriers.²

Multiple barriers contribute to investment stalling, including cost gaps, uncertain demand, and insufficient offtake commitments (International Energy Agency, 2025a). This paper isolates one margin that existing frameworks have not formalized: the timing of bilateral commitment to certified green supply under recognition uncertainty. This uncertainty is especially acute in nascent bilateral markets without standardized contracts or benchmarks. When recognition rules are still forming, waiting has value because early commitment can lock a relationship into standards that may shift. Buyers can continue to rely on incumbent fossil supply, so postponing commitment remains feasible while the recognition environment settles. The flexibility of that incumbent fallback shapes the value of waiting. What determines when commitment happens, and how does incumbent flexibility interact with recognition uncertainty?

I develop a bilateral timing framework in which the buyer commits to certified green supply under recognition uncertainty while retaining an incumbent outside option. Most formal analyses treat the importer’s problem as a static allocation across sources. I instead take commitment timing as the central decision margin. I isolate a dimension that existing frameworks overlook: uncertainty about whether certified output will be recognized as green under importing rules. I call this object eligibility variance. The degree to which this uncertainty has been resolved is certification clarity.

The analytical framework builds on four strands of literature. First, I adopt a real-options view of investment under uncertainty (Arrow and Fisher, 1974; Henry, 1974; Bernanke, 1983; McDonald and Siegel, 1986; Dixit and Pindyck, 1994; Pindyck, 1993; Abel et al., 1996; Bloom, 2009). This literature typically does not give an explicit role to a persistent outside option during a transition. I embed the real-options logic in a bilateral relationship in which the importer retains an incumbent fallback whose flexibility is a primitive of the commitment problem.

Second, I draw on the property-rights and hold-up literature (Grossman and Hart, 1986; Hart and Moore, 1990; Rogerson, 1992; Nunn, 2007). Joskow (1987) provides foundational empirical evidence that relationship-specific investments determine contract duration in bilateral energy supply, using coal contracts between suppliers and electric utilities. In international trade, Handley and Limaão (2015, 2017) show that policy uncertainty interacts with sunk entry costs to delay entry. Zahur (2025) develops a multi-stage model of LNG contracting with spot trade as an outside option, and Antweiler and

¹The IEA reports that for the first time, potential low-emissions hydrogen production by 2030 based on announced projects has declined, from 49 Mtpa to 37 Mtpa between the 2024 and 2025 reviews (International Energy Agency, 2025a).

²In a CEO sentiment survey of over 70 Hydrogen Council member companies, regulatory clarity ranked among the top three success factors for clean hydrogen projects, above supply reliability and technology track record (Hydrogen Council and McKinsey & Company, 2025).

Schlund (2024) study emerging hydrogen trade under environmental and trade policies. I differ in three respects: I study cross-fuel transition rather than within-fuel contracting; the uncertainty is eligibility-based rather than demand-based; and the model generates a regime boundary under which the effect of incumbent flexibility reverses sign rather than remaining monotonic.

Third, I connect to the literature on environmental standards and certification. Ganslandt and Markusen (2001) show how standards can function as trade barriers even when non-discriminatory, and Fischer and Lyon (2014) study the economics of environmental certification. For credence goods, certification creates value but also uncertainty about what recognition implies (Darby and Karni, 1973; Dulleck and Kerschbamer, 2006). Lizzeri (1999) analyzes certification intermediaries who strategically choose disclosure rules; here the uncertainty arises from regulatory fragmentation across jurisdictions. Harbaugh et al. (2011) show that uncertainty about what a standard means can eliminate firms' incentive to seek certification. I extend this logic to a bilateral investment-timing framework where eligibility uncertainty governs the option value of waiting.

Fourth, I build on the literature analyzing investment under carbon policy uncertainty (Hassett and Metcalf, 1999; Blyth et al., 2007; Kellogg, 2014; Boomsma et al., 2012; Nishide and Nomi, 2009; Nishide and Yagi, 2016; Dalby et al., 2018; Flora and Tankov, 2023). Fuchs et al. (2024) show that carbon price uncertainty affects decarbonization investment with magnitude comparable to the effect of price levels, but that channel operates through carbon charges rather than through product eligibility. Delfgaauw and Swank (2025) analyze commitment to green subsidies under domestic redistributive concerns rather than bilateral certification uncertainty, and Besley and Persson (2023) study green transitions under limited political commitment more broadly. Recent work documents that hydrogen certification schemes diverge in lifecycle boundaries, emissions thresholds, and verification requirements (White et al., 2021; Gale et al., 2024; Hussain and Syron, 2026), and that these mismatches materially affect cost competitiveness (White et al., 2025). What remains underdeveloped is a theory of how recognition uncertainty enters the timing of irreversible commitment and how it interacts with incumbent flexibility. This paper provides that framework.

I show that certification clarity governs whether incumbent-market flexibility delays or accelerates bilateral commitment to certified green supply. When recognition rules remain unsettled, flexibility sustains delay because the importer can continue on fossil supply while the certification environment resolves, and greater flexibility raises the value of waiting. When recognition is clear, the same flexibility reduces the downside risk of the green relationship and brings commitment forward. The sign reverses at a characterizable boundary in eligibility variance. This sign reversal recasts the bridge-versus-destination debate. The incumbent fossil market is not inherently one or the other, and the same bilateral relationship can obstruct green commitment under unresolved recognition and support it once certification is clear. I also show that reducing eligibility variance can matter more for commitment timing than additional carbon-price support for certified goods with near-zero embedded emissions, because certification bounds the charges the green good faces under recognition and leaves eligibility uncertainty as the dominant friction. In early-stage bilateral markets, mutual recognition and methodology harmonization can therefore accelerate commitment more effectively than additional carbon-price support alone.

The problem structure extends beyond green hydrogen. Certified green commodities are credence goods whose value depends on attributes that buyers cannot directly

verify and that must be certified or mutually recognized.³ Examples include verified carbon offsets, sustainably certified agricultural commodities, and goods subject to mutual recognition under trade agreements. The distinctive feature is that eligibility uncertainty, whether the good will be recognized as having the claimed attribute, is distinct from price uncertainty, the value conditional on recognition.

Commitment timing is governed by a threshold in the exporter’s alignment effort, the minimum alignment at which the importer finds it worthwhile to forgo incumbent supply and commit to the green relationship. The threshold and its comparative statics do not depend on functional form. The framework also separates this commitment threshold from the exporter’s private ability to finance the required alignment. Delay can therefore persist even when earlier commitment would be jointly beneficial under feasible transfers, because the exporter’s surplus from the relationship may fall short of the cost of reaching the threshold.

I illustrate the model with a calibration to Australia–Japan energy trade, using LNG as the incumbent market and green hydrogen as the certified supply chain. A clarity index constructed from observable institutional developments, such as ISO standard publications, Japanese legislative and regulatory milestones, and bilateral recognition-related developments, traces the evolution of eligibility variance from 2019 to 2024. The baseline index includes both hard institutional events and soft bilateral signals; robustness checks exclude the latter. Under the baseline calibration, the bilateral relationship lies near the regime boundary, and the continued absence of hydrogen FIDs is consistent with alignment effort remaining below the commitment threshold. Descriptive evidence from the global hydrogen project pipeline points in the same direction. Export-oriented electrolysis projects, which are most exposed to eligibility uncertainty, reach final investment decisions at markedly lower rates than domestic projects and than export projects with more established certification pathways.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 sets up the bilateral commitment problem. Section 4 develops the theory of commitment timing, showing how certification clarity governs whether incumbent flexibility delays or accelerates commitment and when eligibility variance has more leverage than carbon price levels on the timing margin. Section 5 analyzes how certification clarity lowers the commitment threshold, shows that incumbent flexibility amplifies the effect of clarity reform, and characterizes instrument interactions. Section 6 calibrates the model to Australia–Japan energy trade. Section 7 presents descriptive evidence from the global hydrogen project pipeline. Section 8 concludes.

2 Background

Cross-border green hydrogen projects have proliferated in announcements but stalled in execution. Certification regimes remain fragmented, and incumbent LNG markets have become more flexible. Together, project stagnation, regulatory uncertainty, and growing incumbent-market flexibility define the setting for the analysis.

³The credence-good concept originates with [Darby and Karni \(1973\)](#) and has been developed in models where experts privately observe a customer’s needs and can overtreat or undertreat when outcomes are not verifiable (e.g., [Dulleck and Kerschbamer, 2006](#); [Dulleck et al., 2011](#)). In the present setting, the credence attribute is the lifecycle carbon intensity that the exporter claims and that the importer’s regime must recognize. The friction is not hidden treatment choice by an expert, but public uncertainty about how this certified attribute will be treated over time in the importing country’s policy regime.

2.1 Cross-border green hydrogen projects

Australia, Chile, Saudi Arabia, and several African countries have positioned themselves as exporters of green hydrogen and its derivatives, particularly ammonia, to energy-importing regions in Europe and East Asia. Japan, South Korea, and Germany have announced import targets and signed memoranda of understanding with prospective suppliers.

Announced ambition has far outpaced realized investment, despite sharp cost declines and substantial policy commitments. Renewable electricity costs have fallen by roughly 90 percent since 2010 ([International Renewable Energy Agency, 2024](#)). The United States' Inflation Reduction Act provides production tax credits of up to US\$3 per kilogram,⁴ the European Union has approved over €18 billion in public funding through four waves of Important Projects of Common European Interest ([European Commission, 2024](#)), and Japan and South Korea have announced import targets backed by publicly supported competitive offtake ([International Energy Agency, 2025a](#)). Yet only about 9 percent of announced low-emissions hydrogen capacity has reached final investment decision, the 2030 project pipeline contracted for the first time since tracking began, and offtake agreements declined in both volume and commitment firmness in 2024 ([International Energy Agency, 2025a](#)).

Among the barriers the IEA identifies ([International Energy Agency, 2025a](#)), certification uncertainty is particularly acute for export-oriented projects. Bankability depends not only on expected levelized costs but also on whether the green hydrogen will be recognized as green in the destination market and whether that recognition will remain stable over the project's lifetime ([Lee and Saygin, 2023](#)).

2.2 Certification and carbon border adjustments

International standards for hydrogen fuel quality and greenhouse gas accounting have been developed by the International Organization for Standardization and the International Partnership for Hydrogen and Fuel Cells in the Economy (IPHE).⁵ These standards provide technical foundations, but adoption and mutual recognition across jurisdictions remain incomplete.⁶ Differences in policy approach, environmental thresholds, and jurisdictional scope can matter even when standards appear similar on paper ([McDermott et al., 2008](#)).

The European Union's Carbon Border Adjustment Mechanism (CBAM) illustrates how third-party regulatory regimes shape methodology expectations beyond their own jurisdictions ([Clausing and Wolfram, 2023](#)). The CBAM requires importers of certain

⁴The One Big Beautiful Bill Act of 2025 accelerated the phase-out of these credits; facilities must begin construction by December 31, 2027, to remain eligible.

⁵ISO 14687:2025 specifies minimum quality characteristics of hydrogen fuel ([International Organization for Standardization, 2025](#)). ISO/TS 19870:2023 provides a methodology for determining greenhouse gas emissions associated with hydrogen production, conditioning, and transport ([International Organization for Standardization, 2023](#)). The methodology provides a common calculation framework but does not set a threshold to qualify hydrogen as low-emission; thresholds are determined by national legislation ([International Organization for Standardization, 2024](#)). The IPHE has published a working paper proposing a methodology for calculating greenhouse gas emissions across hydrogen value chains ([International Partnership for Hydrogen and Fuel Cells in the Economy, 2023](#)).

⁶A 2024 survey commissioned by Japan's Agency for Natural Resources and Energy found that no country had established a finalized methodology for mutual recognition of hydrogen certification, and that Japan itself lacked a domestic hydrogen certification regime at the time of the survey ([Japan Hydrogen Association, 2024](#)).

goods, including hydrogen, to purchase certificates corresponding to embedded carbon, using either EU-provided default values or actual emissions data verified under EU-approved methodologies (European Commission, 2023). The certificate price is linked to the EU Emissions Trading System allowance price. Measurement, reporting, and verification are central to CBAM implementation (OECD, 2025), and interoperability across certification schemes is crucial for cross-border trade (International Energy Agency, 2025a). The CBAM is discussed here as a general archetype of eligibility uncertainty rather than as a direct input into the Australia–Japan calibration in Section 6.

Recent bilateral and domestic initiatives address eligibility uncertainty directly. In October 2024, Indonesia and Japan signed a Mutual Recognition Arrangement confirming equivalence between Japan’s Joint Crediting Mechanism and Indonesia’s greenhouse gas certification (Ministry of the Environment, Government of Japan, 2024). Japan’s Hydrogen Society Promotion Act (2024) establishes a carbon intensity threshold for low-carbon hydrogen, requires ministerial approval of qualifying business plans, and provides price-gap support for approved projects (Agency for Natural Resources and Energy, Ministry of Economy, Trade and Industry, Government of Japan, 2024). Projects that do not meet the threshold, or whose certification cannot be verified under the Act’s criteria, are excluded.

The distinction between soft bilateral agreements and binding recognition reforms matters for investment timing. Memoranda of understanding, feasibility partnerships, and exploratory offtake agreements signal interest and preserve future optionality, but they need not resolve the recognition uncertainty that governs bankability. A corridor may accumulate substantial soft bilateral activity while the underlying eligibility environment remains unsettled. Binding recognition instruments, such as mutual recognition agreements, harmonized methodology adoption, and legislated certification thresholds, directly reduce eligibility uncertainty by settling how green attributes will be assessed and whether that assessment will persist. The calibration in Section 6 distinguishes between these two categories.

For the model, the key implication is that green hydrogen exporters face uncertainty not about the size of any carbon charge, but about whether their certified output will be recognized as low-emission under the importing jurisdiction’s rules and whether that recognition will persist over the project’s lifetime.

2.3 Flexibility in LNG markets and the ‘Bridge’ vs. ‘Destination’ debate

For Japan and other East Asian importers considering green hydrogen, the relevant incumbent market is liquefied natural gas. LNG trade historically relied on long-term contracts with rigid destination clauses, take-or-pay obligations, and oil-indexed pricing. Over the past two decades, contracts have become substantially more flexible (Hartley, 2015). Destination clauses have been relaxed, volume bands widened, and price indexation diversified toward spot gas. The Japan Fair Trade Commission documented this shift in a 2017 report on destination restrictions (Japan Fair Trade Commission, 2017), and the International Group of Liquefied Natural Gas Importers reports that the share of spot and short-term trade grew from under 20 percent in 2010 to over 35 percent by 2023 (GIIGNL, 2024).

This flexibility has direct implications for importers. Destination-flexible contracts and liquid spot markets allow importers to adjust volumes, arbitrage regional price differences, and reduce supplier dependence. An importer can continue to source LNG while

certification uncertainty remains unresolved and commit to the green supply chain after eligibility uncertainty falls.

The Australia–Japan energy relationship exemplifies this dynamic. Japan is the world’s largest LNG importer and Australia its largest supplier. Japanese utilities hold long-term contracts with Australian LNG projects that include destination flexibility and volume adjustment provisions. At the same time, Australian states have announced green hydrogen export projects targeting Japanese buyers, and Japanese utilities have signed preliminary agreements to explore imports. The central question is when those preliminary agreements will convert into binding commitments backed by final investment decisions, and what role LNG flexibility plays in that timing.

The role of LNG in decarbonization remains contested. Proponents emphasize that coal-to-gas switching has delivered measurable emissions reductions, with the IEA estimating over 500 million tons of CO₂ avoided between 2010 and 2018 ([International Energy Agency, 2019](#)). Critics emphasize carbon lock-in, stranded-asset risk ([Kemfert et al., 2022](#)), and methane leakage ([Howarth, 2014](#)). The debate is structured by competing political framings rather than by convergent empirical evidence ([Janzwood and Millar, 2022](#)). This paper does not take a position on that debate. It derives conditions under which LNG flexibility functions as a bridge that insures the transition and conditions under which the same flexibility sustains delay and makes LNG a destination.

3 Model

The model formalizes the commitment-timing problem as a bilateral relationship in which the importer chooses when to commit to certified green supply while the recognition environment remains unsettled and an incumbent fossil fallback remains available. Delayed commitment is represented through the option value of waiting rather than through an explicit model of the full post-waiting sequence. This reduced-form choice keeps the analysis focused on the immediate commitment threshold and its comparative statics with respect to recognition uncertainty and incumbent flexibility.

3.1 Timing and payoffs

At $t = 0$, the exporter chooses alignment effort $e \in [0, 1]$ at convex cost $C_E(e)$, with $C_E(0) = 0$, $C'_E > 0$, and $C''_E > 0$. Alignment effort is a reduced-form exporter-side choice that summarizes actions affecting the recognition environment.⁷ The importer observes e and then decides whether to make an immediate binding commitment to the green supply chain or to wait.⁸

If the importer commits at $t = 0$, it enters an irreversible relationship with the exporter and forgoes access to the incumbent market. If the importer waits, it continues trading

⁷In practice, alignment effort aggregates exporter-controlled investment (certification infrastructure, supply-chain transparency), investment in alignment with prevailing and emerging certification standards including participation in standard-setting and technical validation processes, and responses to importer-side policy signals. The scalar representation captures the net effect of these heterogeneous activities on the recognition environment. Appendix B shows that the threshold structure is robust to alternative timing assumptions, including simultaneous moves and noisy signals.

⁸The model treats commitment as a binary and irreversible move from non-commitment to binding green commitment. In practice, this transition is staged through non-binding memoranda, preliminary offtake agreements, and front-end engineering before a final investment decision, and importers may maintain mixed portfolios during the transition. The binary structure isolates the timing margin at which the relationship becomes binding.

in the incumbent market and retains the option to commit later once uncertainty about certification and border treatment has resolved. The state of the incumbent market is summarized by a flexibility index $\phi \in [0, \bar{\phi}]$, which captures destination rights, volume adjustment provisions, supplier diversification, and spot market access.

3.2 Recognition as a threshold-crossing problem

The green good is a credence good whose value depends on attributes (lifecycle carbon intensity, production pathway, and the additionality of renewable inputs) that the importer cannot directly verify and that must be certified under the importing jurisdiction's rules. Recognition is the event that the importer's regime accepts the exporter's certification and treats the good as zero-carbon for border adjustment purposes.

Definition 1 (Recognition gap). *Let G denote the recognition gap, the difference between the importer's recognition threshold and the exporter's compliance score. Recognition occurs if and only if $G \leq 0$. Lower values of G indicate better alignment with the importer's requirements.*

Uncertainty enters through both sides of G . The exporter's measured compliance depends on evolving lifecycle-accounting methodologies, while the importer's effective standard depends on regulatory discretion and the stringency of border rules. I impose a location-scale structure on the distribution of G conditional on observables.

Assumption 1 (Location-scale representation). *Conditional on alignment effort e , exporter credibility $\kappa_E \geq 0$, and institutional clarity $s \in [0, \bar{s})$ with $\bar{s} \leq 1$, the recognition gap satisfies*

$$G = m(e, \kappa_E, s) + \sqrt{u(e, \kappa_E, s)} \varepsilon, \quad (1)$$

where m is the conditional mean, $u \geq \underline{u} > 0$ is the conditional variance, and ε has mean zero, unit variance, continuous distribution function F , and strictly positive density f on $\mathbb{R}$.

The parameter κ_E captures inconsistency in the exporter's certification record; higher κ_E reflects past behavior or documentation practices that raise questions about the reliability of certified attributes. The parameter s captures institutional clarity, with higher s corresponding to more stable and mutually recognized certification methodologies. The restriction $s < \bar{s}$ and $u \geq \underline{u} > 0$ ensures that some residual uncertainty always remains.⁹

Assumption 2 (Moment shifts from effort). *The conditional mean m and variance u satisfy $m_e < 0$ and $u_e < 0$: alignment effort reduces both the expected gap and its dispersion.*

Higher alignment effort shifts the distribution of G toward recognition through two channels: the mean channel improves the exporter's compliance score, and the variance channel compresses the range of possible assessment outcomes. The distinction between alignment effort e and institutional clarity s reflects different loci of control. Alignment effort captures what the exporter does: investing in process adaptation, documentation, and conformity with certification requirements. Institutional clarity captures the state of the recognition regime: how settled the methodology is, whether mutual recognition

⁹In the calibration of Section 6, I construct an empirical proxy for s from observable policy events and map it to variance through the relation $u = \underline{u} + (\bar{u} - \underline{u})(1 - s/\bar{s})$, which satisfies the lower bound.

is in place, and how predictable border treatment will be over the project lifetime. In practice, these interact: exporter-side compliance investments contribute to the institutional knowledge base, and clearer regimes raise the return to alignment effort. The model separates the two by assigning e to the exporter's decision problem and s to the institutional environment that the exporter takes as given. A later section asks how the recognition environment is itself determined.

Assumption 3 (Variance shifts from clarity and credibility). *The conditional variance satisfies $u_s < 0$ and $u_{\kappa_E} > 0$: higher institutional clarity reduces eligibility variance, and higher exporter inconsistency raises it.*

Definition 2 (Recognition probability). *The probability that the importer's regime recognizes the exporter's certification is*

$$p(e, \kappa_E, s) \equiv \Pr(G \leq 0 \mid e, \kappa_E, s) = F\left(-\frac{m(e, \kappa_E, s)}{\sqrt{u(e, \kappa_E, s)}}\right). \quad (2)$$

The recognition probability depends on the standardized distance $d \equiv -m/\sqrt{u}$. I impose that institutional clarity and exporter inconsistency shift d monotonically.

Assumption 4 (Monotone standardized distance). *The standardized distance $d(e, \kappa_E, s) \equiv -m(e, \kappa_E, s)/\sqrt{u(e, \kappa_E, s)}$ satisfies:*

- (i) $\partial d/\partial s > 0$: *higher institutional clarity improves the importer's standardized assessment of the exporter.*
- (ii) $\partial d/\partial \kappa_E < 0$: *higher exporter inconsistency worsens the standardized assessment.*

Assumption 4 is an economic restriction requiring that the net effect of clarity and inconsistency on the standardized distance has an unambiguous sign.¹⁰ The assumption directly implies $p_s > 0$ and $p_{\kappa_E} < 0$.

3.3 Net gain from early commitment

Let $A(e; \phi, \bar{\tau}, \kappa_E, s)$ denote the *static net gain* from committing at $t = 0$ relative to waiting, holding fixed the option value of delay. This term collects the importer's incremental rents from the green relationship, foregone first-mover rents from delay, baseline risk premia, and expected carbon charges. Let $\Omega(\phi, u)$ denote the *option value of waiting*, which captures how eligibility dispersion affects the value of learning before committing, with the mean component absorbed into A .¹¹

Definition 3 (Net gain from early commitment). *The importer's net gain from committing at $t = 0$ is*

$$\Pi(e; \phi, \bar{\tau}, \kappa_E, s) \equiv A(e; \phi, \bar{\tau}, \kappa_E, s) - \Omega(\phi, u(e, \kappa_E, s)). \quad (3)$$

The importer commits if and only if $\Pi \geq 0$.

¹⁰ A sufficient condition for $\partial d/\partial s > 0$ is $m_s < (m/2u)u_s$. When $m > 0$ and both $m_s < 0$ and $u_s < 0$, this requires that the mean improvement is not fully offset by the variance compression. When $m < 0$, the condition is easier to satisfy because variance compression then reinforces the mean effect. Analogous conditions apply to κ_E .

¹¹ Appendix E.3 derives an option-exercise environment in which the continuation value depends on both m and u . The reduced form $\Omega(\phi, u)$ isolates the dependence on dispersion, while the level effect of the mean enters expected payoffs under commitment.

The model does not solve for the optimal timing of delayed commitment; delayed commitment is represented in the continuation value $\Omega(\phi, u)$, grounded by a stopping-problem construction in Appendix E.4 (Remark 1). The static component A admits a decomposition that separates the carbon-price channels.

Assumption 5 (Static component structure). *The static net gain satisfies*

$$A(e; \phi, \bar{\tau}, \kappa_E, s) = \tilde{A}(e; \phi, \kappa_E, s) + \Delta^F(\bar{\tau}) - \bar{\tau} [p(e, \kappa_E, s) \iota^g + (1 - p(e, \kappa_E, s)) \iota^{\text{def}}], \quad (4)$$

where:

- (i) $\tilde{A}$ collects incremental green rents $F_I^g(e)$, foregone first-mover rents $\delta_I^g(\phi)$, and base-line risk premia $\rho_0(\phi)$, all independent of the carbon price level.
- (ii) $\Delta^F(\bar{\tau})$ is the foregone-rent channel. Higher carbon prices raise the cost of waiting on the fossil outside option, with $\Delta_{\bar{\tau}}^F \geq 0$.
- (iii) The bracketed term is the expected carbon charge on the green good, where ι^g is the assigned intensity under recognition and $\iota^{\text{def}} > \iota^g$ is the default intensity when recognition fails.

The expected carbon charge depends on recognition probability p because certification determines the assigned intensity. Under recognition, $\iota^g \approx 0$; when recognition fails, the default intensity ι^{def} applies. This structure ensures that carbon policy affects timing through the recognition channel rather than through a mechanical accounting identity.¹²

The option value satisfies regularity conditions that arise from the structure of the importer's waiting problem rather than being imposed as primitives.

Assumption 6 (Option value properties). *The function $\Omega : [0, \bar{\phi}] \times [\underline{u}, \bar{u}] \rightarrow \mathbb{R}_+$ is twice continuously differentiable and satisfies:*

- (i) $\Omega_u(\phi, u) > 0$: higher eligibility variance raises the option value of waiting.
- (ii) $\Omega_\phi(\phi, u) \geq 0$: greater incumbent flexibility weakly raises the option value.
- (iii) $\Omega_{\phi u}(\phi, u) > 0$: the marginal value of flexibility is higher when variance is higher.

Condition (i) is the standard convexity property of options. Condition (ii) reflects that a more flexible outside option expands the importer's choice set during the waiting period. Condition (iii), increasing differences in (ϕ, u) , drives the regime-switch result in Section 4.3: when certification is uncertain, flexibility is more valuable because there is more to learn by waiting; when certification is clear, the marginal value of flexibility diminishes.

Remark 1 (Stopping-problem foundation). *The maintained properties in Assumption 6 are grounded by a standard 2-period stopping environment in which the importer chooses between committing now and waiting to observe eligibility resolution. Appendix E.4 provides the derivation and one primitive configuration under which the crossing condition in Assumption 10 holds. These are existence and grounding results; the reduced-form model does not require the specific stopping structure.*

Lemma 1 (Sufficient conditions for increasing differences). *Suppose:*

- (i) *The eligibility state follows a location-scale family: $X = m + \sqrt{u}\varepsilon$ with ε having continuous density $f > 0$ on its support;*

¹²Without the default-intensity term, the carbon-price channel would vanish for certified goods with $\iota^g \approx 0$. The default-intensity structure captures the penalty for non-recognition and generates a non-trivial interaction between carbon policy and eligibility uncertainty.

- (ii) Incumbent flexibility reduces the effective commitment threshold: $K'(\phi) < 0$;
 - (iii) The importer is in the delay region: $m < K(\phi)$.
- Then $\Omega_{\phi u}(\phi, u) > 0$, and Assumption 6(iii) is satisfied.

Proof. See Appendix E.3. □

Increasing differences arise when the importer waits because the expected recognition gap exceeds the commitment threshold, and flexibility lowers that threshold. In this region, higher eligibility variance raises the probability that a favorable realization triggers commitment, and flexibility amplifies this channel.¹³

3.4 Single crossing and threshold existence

The importer's decision reduces to a threshold in alignment effort under a single-crossing condition.

Assumption 7 (Single crossing in alignment effort). *For each $(\phi, \bar{\tau}, \kappa_E, s)$, the function $e \mapsto \Pi(e; \phi, \bar{\tau}, \kappa_E, s)$ is continuously differentiable on $[0, 1]$ and satisfies $\Pi_e(e; \phi, \bar{\tau}, \kappa_E, s) > 0$ for all $e \in (0, 1)$.*

Assumption 7 imposes that the combined effects of alignment effort on static gains and on the option value yield strict monotonicity of Π in e .¹⁴

Assumption 8 (Interior solution). *For each $(\phi, \bar{\tau}, \kappa_E, s)$ in the relevant parameter space, $\Pi(0; \phi, \bar{\tau}, \kappa_E, s) < 0$ and $\Pi(1; \phi, \bar{\tau}, \kappa_E, s) > 0$.*

Assumption 8 ensures that commitment is not optimal at zero effort but is optimal at maximal effort. Together with continuity and single crossing, this guarantees an interior threshold.

4 Commitment threshold and regime classification

Commitment to certified green supply requires the exporter's alignment effort to cross a threshold. The threshold is the minimum alignment at which the importer finds it worthwhile to forgo incumbent fossil supply and commit irreversibly to the green relationship. Whether the exporter can privately afford to supply that alignment is a separate question, because the exporter's surplus from the relationship may fall short of the cost of reaching the threshold. The role of outside-option flexibility is ambiguous: greater flexibility could raise the value of waiting and delay commitment, or reduce the downside risk of the green relationship and accelerate commitment. The direction depends on how resolved eligibility uncertainty is.¹⁵ Formal proofs are collected in Appendix A.

¹³Appendix E.3 provides the full derivation and verifies the conditions hold in the parameterized environment used for calibration.

¹⁴A sufficient condition is $\tilde{A}_e + \bar{\tau}(\iota^{\text{def}} - \iota^g)p_e - \Omega_u u_e > 0$. Since $\Omega_u > 0$ and $u_e < 0$, the term $-\Omega_u u_e > 0$ always contributes positively to Π_e . The condition requires that this contribution, together with $\tilde{A}_e > 0$, outweighs any adverse effect when $p_e < 0$.

¹⁵The linear-quadratic specialization in Appendix F provides closed-form expressions that sharpen magnitudes but do not alter signs.

4.1 Threshold existence and comparative statics

Under Assumptions 7 and 8, a unique threshold exists for each parameter configuration. The threshold and its comparative statics do not depend on functional form.

Theorem 1 (Threshold existence and uniqueness). *Under Assumptions 7 and 8, for each $(\phi, \bar{\tau}, \kappa_E, s)$ there exists a unique threshold $e_g^*(\phi, \bar{\tau}, \kappa_E, s) \in (0, 1)$ such that the importer commits at $t = 0$ if and only if $e \geq e_g^*(\phi, \bar{\tau}, \kappa_E, s)$.*

The threshold e_g^* is the minimum alignment effort that induces the importer to commit. It summarizes the balance between the importer's reasons to wait (baseline risk premia, eligibility variance, option value tied to outside-option flexibility, and foregone rents from delay) and the exporter's ability to close this gap through alignment effort. The structure is analogous to a gap-over-slope formula, where the numerator collects the primitives that favor delay and the denominator collects the effectiveness of effort at overcoming them.¹⁶

Theorem 2 (General comparative statics). *Fix $(\phi, \bar{\tau}, \kappa_E, s)$. For any scalar parameter x that enters Π smoothly, the implicit threshold $e_g^*(x; \phi, \bar{\tau}, \kappa_E, s)$ satisfies*

$$\frac{\partial e_g^*}{\partial x} = -\frac{\Pi_x(e_g^*; \phi, \bar{\tau}, \kappa_E, s)}{\Pi_e(e_g^*; \phi, \bar{\tau}, \kappa_E, s)}. \quad (5)$$

Since $\Pi_e > 0$, the sign of $\partial e_g^*/\partial x$ is opposite to the sign of Π_x .

Theorem 2 provides the backbone for all comparative statics results. Since $\Pi_e > 0$, parameters that raise the net gain from commitment ($\Pi_x > 0$) lower the threshold, and parameters that lower it raise the threshold.

Corollary 3 (Variance, clarity, and consistency effects). *Under Assumptions 3, 4, and 6:*

- (i) *Under the reduced-form baseline in which eligibility variance enters Π only through Ω , $\partial e_g^*/\partial u = \Omega_u/\Pi_e > 0$: higher eligibility variance raises the threshold. When A inherits a u -channel through recognition probability, the derivative generalizes to $(\Omega_u - A_u)/\Pi_e$, which is the operative form in Section 4.4; the sign requires $\Omega_u > A_u$ (Proposition 3).*
- (ii) *$\partial e_g^*/\partial s < 0$: higher institutional clarity lowers the threshold.*
- (iii) *$\partial e_g^*/\partial \kappa_E > 0$: higher exporter inconsistency raises the threshold.*

Certification clarity lowers the threshold through both the recognition-probability channel and the option-value channel. Reducing eligibility variance through mutual recognition, methodology harmonization, or pre-certification mechanisms compresses the option value of waiting. Exporter consistency operates through both channels: a more consistent exporter faces a better-positioned distribution and lower variance, both of which lower the threshold.

Corollary 4 (Technological optimism and delay). *Let $\xi \geq 0$ index expected green technology cost drift. Suppose $\Omega_\xi > 0$ (waiting is more valuable when future costs are expected to fall) and $\delta_{I,\xi}^g < 0$ (first-mover rents shrink when all parties expect cost declines). Then $\partial e_g^*/\partial \xi > 0$: stronger expectations of future cost reductions raise the threshold and delay commitment.*

¹⁶Appendix F makes this explicit. Under simple linear-quadratic functional forms, $e_g^* = N(\phi)/D$, where N collects the reasons to wait and $D > 0$ measures the closing power of alignment. Under the interaction specification used for the bridge/destination cutoff, the denominator becomes $D(\phi)$; see (44).

Corollary 4 formalizes the intuition that favorable cost projections do not automatically accelerate green investment. When costs are expected to fall, the option to wait becomes more valuable because the importer can access cheaper technology by delaying. First-mover rents compress because early commitment confers less advantage when everyone expects the frontier to improve. On the timing margin, technological optimism works against early commitment unless paired with variance reduction or rent-sharing arrangements that offset the option-value channel.¹⁷

Mean versus variance channels for recognition. Alignment effort affects the recognition probability through two channels. The mean channel shifts the expected gap toward compliance. The variance channel compresses the dispersion of assessment outcomes. The relative importance of these channels depends on the position of the distribution relative to the recognition threshold. Appendix C (Proposition 6) states the formal decomposition and the sign logic used in this section.

The mean channel always raises recognition probability: shifting the expected gap toward compliance ($m_e < 0$) moves probability mass toward the recognition region regardless of where the distribution sits. The variance channel depends on whether the exporter is expected to be above or below the recognition threshold.

In early-stage certified-goods export corridors, a natural case is $m > 0$, meaning that the exporter is expected to fall short and recognition requires a favorable realization. In this region, variance compression alone can reduce recognition probability by shrinking the left tail that delivers lucky recognition. Alignment effort can still raise total recognition probability if the mean shift dominates the adverse variance-channel effect, a condition that holds when effort moves the distribution sufficiently toward compliance. This decomposition is distinct from the effect of variance on the option value of waiting, which depends on u through the value of learning before committing rather than through recognition probability directly.

4.2 Exporter equilibrium and the implementation wedge

The threshold e_g^* is the minimum alignment at which the importer finds commitment worthwhile. Whether the exporter can privately supply that alignment is a separate question. Let $B_E > 0$ denote the exporter's net surplus from an established green relationship. The exporter finds it privately profitable to induce commitment by supplying effort $e = e_g^*$ if and only if $B_E \geq C_E(e_g^*)$. When this inducement condition fails, the exporter supplies $e = 0$ and no relationship forms.¹⁸

The comparative statics of Theorems 1 and 2 translate directly into comparative statics for equilibrium investment. Parameters that lower the threshold make the inducement condition $B_E \geq C_E(e_g^*)$ easier to satisfy and expand the set of relationships that form.

The baseline game abstracts from side payments and cost-sharing instruments. Under that restriction, delay can be inefficient.

¹⁷The result does not imply that falling costs are bad for green investment in general. On the *timing margin*, expected cost declines amplify the value of waiting and can delay commitment even as leveled costs fall.

¹⁸This characterization assumes that effort above the threshold does not yield additional surplus for the exporter. If B_E increases in e above the threshold, the exporter may supply $e > e_g^*$. The threshold e_g^* remains the critical object for whether any commitment occurs.

Proposition 1 (Delay with positive joint gains under transfers). *There exist parameter values and an effort level $\tilde{e} > e_g^*(\phi)$ such that: (i) the importer strictly prefers commitment at $\tilde{e}$, that is $\Pi(\tilde{e}; \phi) > 0$, (ii) the exporter's private inducement condition fails at $\tilde{e}$, so $B_E < C_E(\tilde{e})$ and the equilibrium exhibits delay absent transfers, and (iii) commitment at $e = \tilde{e}$ is jointly beneficial in the sense that there exists a transfer z such that both parties are weakly better off relative to delay.*

Earlier commitment can be jointly beneficial under feasible transfers or co-financing even when it does not arise in equilibrium. Appendix D provides the proof and discusses interpretations.

4.3 Regime switch: bridge versus destination

The comparative statics of Section 4.1 show how individual parameters shift the threshold level but do not settle the direction of the outside-option flexibility effect. Greater flexibility could delay commitment by raising the value of waiting, or accelerate it by reducing the downside risk of the green relationship. Which effect dominates depends on the level of eligibility variance.¹⁹

4.3.1 The sign condition

From Theorem 2, the effect of outside-option flexibility on the threshold is

$$\frac{\partial e_g^*}{\partial \phi} = -\frac{\Pi_\phi(e_g^*; \phi, u)}{\Pi_e(e_g^*; \phi, u)}. \quad (6)$$

Since $\Pi_e > 0$ by Assumption 7, the sign of $\partial e_g^*/\partial \phi$ is opposite to the sign of Π_ϕ .

Definition 4 (Bridge and destination regimes). *The incumbent outside option is in the bridge regime if $\partial e_g^*/\partial \phi < 0$, so that higher flexibility lowers the threshold and accelerates commitment. It is in the destination regime if $\partial e_g^*/\partial \phi > 0$, so that higher flexibility raises the threshold and delays commitment.*

Using the structure $\Pi = A - \Omega(\phi, u)$ from (3), differentiation with respect to ϕ yields

$$\Pi_\phi = A_\phi - \Omega_\phi(\phi, u), \quad (7)$$

where A_ϕ collects the effects of flexibility on static payoffs and Ω_ϕ is the effect of flexibility on the option value of waiting.²⁰ The static effect A_ϕ reflects two opposing forces: flexibility can reduce the baseline risk premium (bridge-favoring) and can reduce foregone rents from delay (destination-favoring). Under the structure in Assumption 5, the u -dependence in A operates through recognition probability, which does not interact with ϕ .

The regime classification depends on whether static effects or option-value effects dominate. If $A_\phi > \Omega_\phi$, the incumbent is a bridge; if $A_\phi < \Omega_\phi$, it is a destination.

¹⁹Appendix F provides a linear-quadratic specialization that yields a closed-form cutoff.

²⁰Appendix E discusses economic primitives for these channels.

4.3.2 Regime switch under increasing differences

The option-value component Ω_ϕ depends on eligibility variance u through the increasing-differences condition in Assumption 6(iii). When u is high, flexibility is more valuable because there is more to learn by waiting. This dependence can generate a regime switch as u varies.

Assumption 9 (Static flexibility effect independent of certification state). *The static component A_ϕ does not depend on eligibility variance u or on institutional clarity s :*

$$A_{\phi u} = 0, \quad A_{\phi s} = 0.$$

Assumption 9 is satisfied under Assumption 5: the static payoff components depend on the flexibility of incumbent contracts rather than on the state of green certification. The components of $\tilde{A}$ that depend on ϕ —the baseline risk premium $\rho_0(\phi)$ and foregone first-mover rents $\delta_I^g(\phi)$ —reflect LNG contract structure, not the variance or clarity of hydrogen certification. The carbon-charge term depends on recognition probability $p(e, \kappa_E, s)$, which does not depend on ϕ , so its ϕ -derivative is zero regardless of s or u .²¹

Assumption 10 (Limit conditions). *There exist $\underline{u}, \bar{u}$ with $0 < \underline{u} < \bar{u}$ such that:*

- (i) $A_\phi - \Omega_\phi(\phi, \underline{u}) > 0$: *at low variance, static effects dominate.*
- (ii) $A_\phi - \Omega_\phi(\phi, \bar{u}) < 0$: *at high variance, option-value effects dominate.*

Assumption 10 ensures that both regimes are attainable: when certification is nearly certain, flexibility primarily provides insurance; when certification is highly uncertain, flexibility primarily raises the value of waiting.

Theorem 5 (Regime switch under increasing differences). *Under Assumptions 6, 9, and 10, for each $\phi \in (0, \bar{\phi})$ there exists a unique cutoff $u^*(\phi) \in (\underline{u}, \bar{u})$ such that:*

- (i) *If $u < u^*(\phi)$, then $\partial e_g^*/\partial \phi < 0$ and the incumbent is in the bridge regime.*
- (ii) *If $u > u^*(\phi)$, then $\partial e_g^*/\partial \phi > 0$ and the incumbent is in the destination regime.*

The result follows because $h(u) \equiv A_\phi - \Omega_\phi(\phi, u)$ is strictly decreasing in u (by $\Omega_{\phi u} > 0$) and changes sign between $\underline{u}$ and $\bar{u}$ (by Assumption 10), so there is a unique root $u^*(\phi)$. The classification is general given any Ω satisfying Assumption 6; the crossing condition in Assumption 10 ensures both regimes are attainable and is grounded by one primitive construction from a standard stopping problem in Appendix E.4 (Lemma 2). Under the linear-quadratic specialization, the cutoff admits a closed-form expression (Corollary 7).

The economic logic of the sign switch has three parts. When recognition is uncertain, outside-option flexibility raises the value of waiting because the importer can continue on incumbent supply while the recognition environment resolves. Flexibility substitutes for commitment. When recognition is clear, flexibility primarily insures the downside of the green relationship: if the relationship underperforms, the importer can fall back without catastrophic loss. Flexibility complements commitment. The regime boundary $u^*(\phi)$ is the point at which the insurance value of flexibility overtakes its waiting value.

The bridge/destination classification is not intrinsic to the fuel or contract structure. The same LNG portfolio can function as a bridge when certification regimes are clear and as a destination when certification is uncertain.

²¹The condition $A_{\phi u} = 0$ can be relaxed at the cost of additional conditions on how A_ϕ varies with u . The key requirement is that A_ϕ not increase in u faster than Ω_ϕ , which would prevent a sign change. For $A_{\phi s}$, the condition $A_{\phi s} \geq 0$ would strengthen the results in Section 5; the restriction $A_{\phi s} = 0$ is the conservative baseline.

Eligibility variance enters the threshold problem twice. Higher u raises the reasons to wait (directly through the option value) and shifts the incumbent toward the destination regime (indirectly through the sign condition). Instruments that reduce eligibility variance, such as mutual recognition or methodology harmonization, therefore operate on both margins simultaneously: they lower the threshold level and shift the incumbent toward bridge.

The result also provides a framework for incumbent contract design. Features that tie flexibility to green milestones, such as volume relief contingent on certification progress, reduce the independent value of the outside option and shift the relationship toward bridge.²² Unconditional flexibility unlinked to green progress makes destination more likely. Assessments of whether LNG crowds out green hydrogen must condition on the state of certification, because the same contract structure can have opposite effects depending on where the relationship sits relative to u^* .

The connection to the hold-up literature clarifies the role of increasing differences. When $\Omega_\phi = 0$, flexibility enters only through static payoffs and the regime classification is unambiguous. If $A_\phi < 0$ (foregone-rent reduction dominates risk-premium reduction), flexibility always raises the commitment threshold. The logic parallels the property-rights literature: a stronger outside option reduces the net value of relationship-specific commitment (Grossman and Hart, 1986; Hart and Moore, 1990). The bridge regime requires $\Omega_{\phi u} > 0$, which arises when the value of waiting depends on both the flexibility of the outside option and the uncertainty about certification outcomes. This interaction is absent from models that treat eligibility as fixed or that omit real-options considerations from bilateral investment. Even in the bridge regime, however, commitment may not occur if the exporter cannot privately finance the required alignment effort. The implementation wedge characterized in Section 4.2 remains binding whenever $B_E < C_E(e_g^*)$, so the regime classification determines the direction of the flexibility effect on timing but does not guarantee that commitment is privately implementable.

4.4 Policy leverage: eligibility variance versus carbon price levels

The regime-switch result establishes that eligibility variance determines whether flexibility helps or hinders commitment. A further question is how eligibility variance compares to carbon price levels as a lever on the timing margin. Both affect the threshold, but for certified goods with near-zero assigned intensity, their magnitudes can differ sharply.

4.4.1 Channel decomposition

Carbon policy affects the timing threshold through two distinct channels: a direct channel operating through expected carbon charges on the green good, and an indirect channel operating through the cost of waiting on the fossil outside option. The structure of certified goods implies that these channels have sharply different magnitudes. From Theorem 2, applying $\partial e_g^*/\partial x = -\Pi_x/\Pi_e$ to the carbon price level $\bar{\tau}$ yields the following decomposition.

²²Appendix F shows how take-or-pay provisions and importer finance modify $u^*(\phi)$. Contracts that tie relief to alignment milestones raise the effective risk-premium reduction χ and lower the effective foregone-rent reduction ξ_δ , both of which raise $u^*(\phi)$ and expand the bridge region.

Proposition 2 (Carbon-price channel decomposition). *Under Assumptions 5 and 7, the sensitivity of the threshold to the carbon price level satisfies*

$$\frac{\partial e_g^*}{\partial \bar{\tau}} = \frac{1}{\Pi_e} \left[\underbrace{p \iota^g + (1-p) \iota^{\text{def}}}_{\text{expected assigned intensity}} - \underbrace{\Delta_{\bar{\tau}}^F}_{\text{foregone-rent channel}} \right]. \quad (8)$$

The first term is the direct effect through expected carbon charges on the green good. The second term is the indirect effect through the fossil outside option.

The sign depends on which channel dominates. If the foregone-rent channel is strong ($\Delta_{\bar{\tau}}^F$ large), higher carbon prices lower the threshold by raising the cost of delay on fossil. If the direct channel dominates, higher carbon prices raise the threshold by increasing expected charges on the green good. For certified goods, the direct channel is attenuated because certification assigns near-zero intensity ($\iota^g \approx 0$) under recognition; the expected assigned intensity is therefore bounded by $(1-p) \iota^{\text{def}} + p \bar{\iota}$ with $\bar{\iota} \approx 0$, and is disciplined by certification design rather than by the carbon-price process.

4.4.2 Variance channel

Eligibility variance affects the timing threshold through an indirect channel operating through recognition probability and hence expected carbon charges, and a direct channel operating through the option value of waiting. From the structure $\Pi = A - \Omega(\phi, u)$,

$$\Pi_u = A_u - \Omega_u, \quad (9)$$

where $A_u = \bar{\tau}(\iota^{\text{def}} - \iota^g)p_u$ and $\Omega_u > 0$. Using $d \equiv -m/\sqrt{u}$, the effect of variance on recognition probability can be written as

$$p_u = -\frac{df(d)}{2u}, \quad (10)$$

which depends on the standardized distance d rather than on m directly. The following condition ensures that p_u and hence A_u are bounded.

Assumption 11 (Regularity of the standardized density). *The density f of the standardized shock ε satisfies $\sup_{x \in \mathbb{R}} f(x) < \infty$ and $\sup_{x \in \mathbb{R}} |x|f(x) < \infty$.*

Assumption 11 is satisfied by all standard location-scale families (normal, logistic, Student- t with $\nu > 1$). Combined with $u \geq \underline{u} > 0$, it ensures that $|p_u|$ and hence A_u are bounded, and implies $f(x) \rightarrow 0$ as $|x| \rightarrow \infty$.²³

Proposition 3 (Eligibility variance effect). *Under Assumptions 1, 6, 7, and 11,*

$$\frac{\partial e_g^*}{\partial u} = \frac{\Omega_u - A_u}{\Pi_e}. \quad (11)$$

The threshold is increasing in eligibility variance ($\partial e_g^*/\partial u > 0$) whenever the option-value channel dominates the expected-charge channel ($\Omega_u > A_u$).

²³To see this, note that $|x|f(x) \leq M$ implies $f(x) \leq M/|x|$ for $|x| > 0$, so $f(x) \rightarrow 0$ as $|x| \rightarrow \infty$.

The condition $\Omega_u > A_u$ holds when expected noncompliance is positive ($m > 0$) and the recognition threshold is sufficiently distant in standardized units. When recognition is far from the threshold in standardized units, the product $|d|f(d)$ remains bounded by Assumption 11 while $f(d)$ is small by tail decay, so A_u is quantitatively muted. The option-value channel Ω_u remains governed by the convexity of the waiting option and can be substantial even when recognition is currently unlikely.²⁴

4.4.3 Dominance condition

The comparison of magnitudes between the two channels determines which policy lever is more effective on the timing margin.

Proposition 4 (Variance dominance condition). *Under Assumptions 5, 6, 7, and 11, suppose $\Omega_u > A_u$ so that $\partial e_g^*/\partial u > 0$. The ratio of sensitivities satisfies*

$$\frac{|\partial e_g^*/\partial u|}{|\partial e_g^*/\partial \bar{\tau}|} = \frac{\Omega_u - A_u}{|\Delta_{\bar{\tau}}^F - p \iota^g - (1-p) \iota^{\text{def}}|}. \quad (12)$$

For certified goods satisfying the intensity bound $\iota^g \leq \bar{\iota}$ with $\bar{\iota} \approx 0$, a sufficient condition for the variance channel to dominate the carbon-price channel by a factor of at least $R > 1$ is

$$\bar{\iota} \leq \frac{\Omega_u - A_u - R |\Delta_{\bar{\tau}}^F - (1-p) \iota^{\text{def}}|}{Rp}. \quad (13)$$

This condition is feasible only if $\Omega_u - A_u > R |\Delta_{\bar{\tau}}^F - (1-p) \iota^{\text{def}}|$ and $p > 0$. The bound is sufficient and conservative, using the triangle inequality and trading sharpness for transparency.

Variance dominance is a structural implication of certified-good design: the numerator measures the net effect of eligibility variance on the threshold, and the denominator measures the net effect of carbon prices through expected charges and foregone rents. For certified goods in the empirically relevant region, A_u is bounded by the tail-decay property while the direct charge channel is attenuated by $\iota^g \approx 0$, leaving the foregone-rent channel $\Delta_{\bar{\tau}}^F$ as the primary carbon-price lever.

Two objects enter the commitment threshold with distinct economic roles. The carbon price level $\bar{\tau}$ governs charges conditional on assigned intensity. The eligibility variance u governs uncertainty about whether favorable treatment will be recognized.

For certified goods with near-zero assigned intensity, the expected carbon charge equals $\bar{\tau}(1-p) \iota^{\text{def}} + \bar{\tau} p \iota^g$, and certification bounds the second term by $\bar{\tau} p \bar{\iota}$. The charge therefore depends primarily on the non-recognition penalty $(1-p) \iota^{\text{def}}$ weighted by the carbon price, with the certified-good component attenuated by the intensity bound. By contrast, eligibility variance enters the option value of waiting directly and, when $\Omega_u > A_u$, without full offset from the expected-charge channel.

Instruments affecting these two objects have different leverage on the timing margin. Instruments that reduce eligibility variance, such as mutual recognition agreements, methodology harmonization, and pre-certification mechanisms, lower the threshold through the option-value channel. Instruments that stabilize or raise carbon prices,

²⁴For the normal distribution, $|x|f(x) = |x|(2\pi)^{-1/2} \exp(-x^2/2)$, which attains its maximum of $(2\pi e)^{-1/2} \approx 0.24$ at $|x| = 1$ and decays for larger $|x|$. The expected-charge channel is correspondingly bounded.

such as carbon contracts for difference (Richstein and Neuhoff, 2022), price floors, and border adjustments, affect the threshold through the expected-charge and foregone-rent channels. For certified goods satisfying the conditions in Proposition 4, reductions in eligibility variance can have larger effects on timing than mean-price instruments.²⁵ The local-slope comparison in Proposition 4 is stated under the model’s maintained normalization of eligibility variance.

The real-options literature on policy uncertainty studies how uncertainty in key policy variables alters investment timing, often through option-value and risk-premium channels, but typically treats product eligibility as given (e.g., Hassett and Metcalf, 1999; Aldy and Armitage, 2020; Blyth et al., 2007). The standards and certification literature identifies regulatory divergence as a barrier to trade but models it as a static cost rather than as uncertainty affecting timing (e.g., Ganslandt and Markusen, 2001; Fischer and Lyon, 2014). Treating eligibility as a separate state variable that enters the option value of waiting unifies these and generates comparative statics that neither captures in isolation.

5 Certification clarity and policy leverage

The analysis so far has taken the importing country’s eligibility regime as exogenous. This section treats institutional clarity s as a policy variable, consistent with the timing game in Sections 3–4.3 but not attempting to model the full political economy of standard-setting. Recall that alignment effort e captures exporter-controlled actions, while institutional clarity s captures the recognition regime that the exporter takes as given. Introducing s as a policy margin completes the bilateral structure: the exporter chooses alignment effort taking the recognition environment as given, while the recognition environment is shaped by importer policy choices and bilateral institutional processes that no individual exporter controls.

5.1 Clarity and the threshold

The institutional clarity parameter $s \in [0, \bar{s})$ captures the stability and mutual recognition of certification methodologies. Higher s reduces eligibility variance u (Assumption 3) and increases the standardized distance $d = -m/\sqrt{u}$ (Assumption 4), which raises recognition probability.²⁶ A separate channel operates through $u_s < 0$, which lowers the option value of waiting.

I collect the monotonicity restrictions on how s affects the components of Π .

Assumption 12 (Clarity effects on payoff components). *The clarity parameter s affects the payoff components as follows:*

- (i) $u_s < 0$: higher clarity reduces eligibility variance (Assumption 3).
- (ii) $p_s > 0$: higher clarity raises recognition probability (implied by Assumption 4).
- (iii) $A_s \geq 0$: higher clarity weakly raises the non-carbon components of static gains, through increased incremental green rents or reduced baseline risk premia.

²⁵This is a partial-equilibrium statement about the certification margin. Carbon pricing remains important for system-wide substitution between green and fossil supply, for the profitability of marginal projects, and for financing.

²⁶For $m > 0$, lowering u alone would make d more negative. The condition $d_s > 0$ therefore requires $m_s < 0$ with sufficient magnitude to offset this effect.

Assumption 12 consolidates the earlier monotonicity restrictions. The option value $\Omega(\phi, u)$ depends on s only through u ; there is no separate direct effect beyond the variance channel.

Proposition 5 (Clarity lowers the threshold). *Under Assumptions 7, 5, 6, and 12, the alignment threshold is strictly decreasing in clarity:*

$$\frac{\partial e_g^*}{\partial s} < 0. \quad (14)$$

The proof (Appendix A) yields an additive decomposition into two reinforcing channels:

$$\frac{\partial e_g^*}{\partial s} = \underbrace{-\frac{A_s}{\Pi_e}}_{\text{static-payoff channel}} + \underbrace{\frac{\Omega_u u_s}{\Pi_e}}_{\text{option-value channel}}. \quad (15)$$

Both terms are negative. The static-payoff channel operates through higher recognition probability ($p_s > 0$) and weakly higher non-carbon static gains ($\tilde{A}_s \geq 0$); the option-value channel operates through variance compression ($u_s < 0$), which reduces the value of waiting.

Different policy instruments load on different channels. Mutual recognition agreements primarily compress eligibility variance and load on the option-value channel. Threshold alignment, such as harmonizing carbon-intensity definitions, shifts the expected compliance position and loads on the static-payoff channel. Methodology harmonization loads on both. The monotonicity of $\partial e_g^*/\partial s < 0$ holds regardless of relative magnitudes, but the decomposition matters for instrument design because different instruments have different cost structures.²⁷

5.2 Flexibility and the value of clarity reform

Proposition 5 establishes that clarity lowers the threshold. A further question is whether the magnitude of this improvement depends on the flexibility of the incumbent market. The regime-switch result in Theorem 5 showed that flexibility and eligibility variance are complements in the option value of waiting: when both are high, the value of delay is largest. Clarity reform works by compressing eligibility variance. If the option value is higher to begin with because the outside option is flexible, there is more to compress, and the reform has greater leverage.

Define the net-gain improvement from a clarity reform $s_0 \rightarrow s_1$ ($> s_0$) at fixed alignment effort e as

$$\Delta\Pi(e, \phi; s_1, s_0) \equiv \Pi(e; \phi, s_1) - \Pi(e; \phi, s_0). \quad (16)$$

Corollary 6 (Flexibility amplifies the net-gain effect of clarity reform). *Under Assumptions 5, 6, 9, and 12, for any fixed alignment effort $e \in [0, 1]$ and any reform that raises clarity from s_0 to $s_1 > s_0$, the net-gain improvement $\Delta\Pi(e, \phi; s_1, s_0)$ is strictly increasing in incumbent flexibility ϕ :*

$$\frac{\partial \Delta\Pi(e, \phi; s_1, s_0)}{\partial \phi} > 0. \quad (17)$$

²⁷In the linear-quadratic calibration, the clarity parameter s enters only through the variance mapping $u = \bar{u}(1 - s)$, so the option-value channel accounts for the entire clarity effect. The static-payoff channel is active in the general model but is not separately parameterized in the LQ specification. This reflects a calibration choice, not a theoretical restriction.

Proof. See Appendix A. □

The amplification operates through the option-value channel, not the static-payoff channel. Clarity reform reduces eligibility variance, which compresses the option value of waiting. When the outside option is more flexible, the option value starts higher because the importer can wait more cheaply, and the compression from variance reduction therefore removes more delay value. The static payoff from commitment does not contribute to the interaction because the components of A_ϕ that flexibility affects, the baseline risk premium and foregone first-mover rents, depend on LNG contract structure and do not vary with the state of hydrogen certification (Assumption 9).²⁸

For any given level of exporter alignment, corridors with more flexible incumbent markets therefore gain more from certification reform. The result holds for any discrete reform $s_0 \rightarrow s_1$, not only for marginal changes, and does not depend on functional form. It is a fixed-effort statement: the reform's effect on the importer's commitment margin is amplified by flexibility, but the exporter's equilibrium response to the reform is not characterized at this level of generality. The linear-quadratic specialization confirms that the amplification extends to the equilibrium threshold ($\partial^2 e_g^* / \partial s \partial \phi < 0$) without additional assumptions (Appendix F), and the calibration in Section 6 quantifies the magnitude.²⁹

5.3 Instrument interaction

When multiple uncertainty-reducing instruments are available, a natural question is whether to deploy them jointly or sequentially. Two classes of instruments are relevant. The first reduces eligibility variance u (mutual recognition, methodology harmonization). The second reduces technical variance σ^2 in the production process (pilot projects, engineering studies, data transparency). The answer depends on the curvature of the option value function.

Appendix F.8 states and proves a local interaction result for these instruments. When technical and eligibility uncertainty are complements in the option value of waiting, reducing one source of uncertainty lowers the marginal reduction in the commitment threshold from reducing the other. Under additivity, the two effects are locally separable.

The interaction result implies a local sequencing rule when policy resources are scarce.³⁰ Under complementarity in the option value ($\Omega_{u\sigma^2} > 0$), reducing one source of uncertainty lowers the marginal benefit of reducing the other. Combined with the variance dominance result from Section 4.4, this suggests prioritizing eligibility clarity over technical de-risking when eligibility variance is the dominant source of delay. Front-loading reductions in u avoids erosion of their marginal effect by subsequent reductions in σ^2 .³¹

²⁸Formally, $\Pi_{\phi s} = A_{\phi s} - \Omega_{\phi u} \cdot u_s$. Under $A_{\phi s} = 0$, the cross-partial reduces to $-\Omega_{\phi u} \cdot u_s > 0$, driven by the same increasing-differences property ($\Omega_{\phi u} > 0$) that generates the regime switch in Theorem 5.

²⁹One can embed the reform problem in a reduced-form importer objective. The corollary provides the key comparative static for that analysis: corridors with more flexible outside options have higher marginal returns to clarity investment. A full political-economy analysis of reform adoption is beyond the scope of this paper.

³⁰The result characterizes marginal effects under the stated local conditions. For large discrete interventions or settings where $\Pi_{ee} \neq 0$, the sequencing implication may be modified by higher-order effects.

³¹The sequencing rule is local and governs the order of marginal moves, not a global corner solution. A planner with a large budget might optimally reduce both variances, but should front-load reductions in the dominant variance.

6 Calibration: Australia–Japan energy trade

The model predicts that eligibility variance raises the commitment threshold, that the effect of incumbent flexibility on timing depends on a regime boundary, and that certification clarity can dominate mean-price instruments when the option-value channel is strong. Whether these predictions are empirically plausible depends on magnitudes that the theory alone cannot determine. This section calibrates the model to Australia–Japan energy trade to discipline three questions: whether the corridor sits near the regime boundary that separates bridge from destination, whether the observed time-series movement in the threshold comes primarily from flexibility or from eligibility clarity, and which instruments have the greatest leverage on commitment timing. The calibration is illustrative rather than estimated; structural parameters are chosen for internal consistency and qualitative pattern-matching.^{32 33}

6.1 Setting and calibration inputs

Japan is the world’s largest LNG importer, and Australia has been its largest supplier, accounting for approximately 42 percent of Japan’s LNG imports in 2023 (GIIGNL, 2025).³⁴ Long-term contracts historically featured rigid destination clauses and take-or-pay obligations, but terms have evolved toward greater flexibility as destination restrictions have been relaxed, volume bands widened, and spot and short-term trade grown (Japan Fair Trade Commission, 2017; Hartley, 2015; Sharples, 2023). Australia and Japan have announced green hydrogen and ammonia projects targeting bilateral trade, but only a small fraction of announced projects have reached FID (International Energy Agency, 2024).³⁵

Outside-option flexibility. The outside-option index ϕ_t is the equally-weighted average of three components of LNG market flexibility—spot share, supplier diversification, and regasification flexibility—each computed from GIIGNL annual reports.³⁶

$$\phi_t = \frac{1}{3} (\varsigma_t + d_t + r_t), \quad (18)$$

³²I adopt the linear-quadratic specialization from Section 4 and Appendix F. The clarity parameter s from Section 5 is mapped to the year-specific clarity index $\tilde{s}_t$ constructed below, with eligibility variance given by $u_t = \bar{u}(1 - \tilde{s}_t)$. In this calibration, the object u corresponds to eligibility variance as defined in Sections 3–5; I maintain σ^2 for technical variance, which enters the instrument-interaction analysis in Section 5.3 but is held fixed in the baseline calibration.

³³Scenario comparisons reflect the model’s maintained normalization of eligibility variance and should not be read as a normalization-invariant empirical dominance test.

³⁴GIIGNL annual reports are organized by market year. The GIIGNL Annual Report edition $t + 1$ covers market year t . The 42 percent figure refers to market year 2023 and is computed from the GIIGNL Annual Report 2024, which covers that year. The GIIGNL Annual Report 2025, covering market year 2024, records a subsequent decline in Australia’s share.

³⁵The Hydrogen Energy Supply Chain project in Victoria’s Latrobe Valley completed a pilot shipment to Kobe in 2022. Bankability of export-oriented projects depends on certification credibility and treatment in the destination market (Lee and Saygin, 2023).

³⁶Year t refers to the market year; GIIGNL Annual Report edition $t + 1$ reports market-year t data. For example, GIIGNL Annual Report 2024 covers calendar year 2023. The three components are computed as follows:

- Spot share: $\varsigma_t = (\text{spot} + \text{short-term volume})_t / (\text{total LNG imports})_t$
- Supplier diversification: $d_t = 1 - \text{HHI}_t$ where $\text{HHI}_t = \sum_j (\text{share}_{j,t})^2$ and $\text{share}_{j,t}$ is the import share from country j
- Regasification flexibility: $r_t = 1 - (\text{throughput}_t / \text{nameplate capacity}_t)$

Data are compiled from GIIGNL Annual Reports 2020–2025, covering market years 2019–2024.

where ς_t is the spot share, d_t is supplier diversification, and r_t is regasification flexibility. Table 1 reports the components and composite index for 2019–2024. The composite index is stable at approximately 0.55 over the sample; the spot share has declined while regasification flexibility has increased.

Eligibility variance. Eligibility variance u_t is constructed from an event-based proxy $\tilde{s}_t$ for the resolution of eligibility uncertainty, not a direct measurement of the recognition gap. The index reflects the cumulative resolution of uncertainty about how green hydrogen will be treated under certification and border rules, and is based on observable policy events with clear dates.³⁷ Eligibility variance is defined as

$$u_t = \bar{u}(1 - \tilde{s}_t), \quad (19)$$

where $\bar{u} = 0.80$ is a normalization representing maximum eligibility variance at zero clarity. Under this construction, eligibility variance falls from 0.68 in 2019 to 0.32 in 2024. The policy scenarios represent comparative statics with respect to exogenous shifts in u , holding fixed the exporter’s behavioral mapping.³⁸

Structural parameters. The remaining parameters are set to produce interior thresholds given the data-driven values of ϕ_t and u_t .³⁹ The default intensity $\iota^{\text{def}} = 0.10$ in normalized units represents the carbon intensity assigned to hydrogen that fails certification; robustness checks vary ι^{def} over $[0.05, 0.15]$ without qualitative change.

In the LQ calibration, the threshold is increasing in eligibility variance throughout the sample. Eligibility variance enters the threshold only through the option-value block; the recognition-probability channel A_u from Proposition 3 is not separately parameterized. The threshold sensitivity is

$$\frac{\partial e_g^*}{\partial u} = \frac{\beta_2 + \xi_{\Omega,1}\phi}{D(\phi)},$$

which requires the numerator $\beta_2 + \xi_{\Omega,1}\phi$ and the single-crossing denominator $D(\phi)$ to be positive. Under the baseline calibration, both conditions hold throughout the sample, with $\partial e_g^*/\partial u$ ranging from 0.514 to 0.517 across years.

In the general model, the condition $\Omega_u > A_u$ additionally requires that the option-value sensitivity exceeds the recognition-probability channel; that condition and its economic interpretation are discussed in Appendix C. The baseline specification assumes

³⁷The baseline clarity in 2019 is set to 0.15, reflecting limited international consensus on hydrogen GHG accounting prior to coordinated methodology development. Clarity increases with subsequent developments: JP Basic Hydrogen Strategy (2020, +0.05), JP Green Growth Strategy (2021, +0.05), IPHE GHG methodology draft (2022, +0.10), ISO 19870 and related harmonization efforts (2023, +0.15), JP GX framework (2024, +0.05), and AUS-JPN hydrogen MoU (2024, +0.05). These increments yield cumulative levels $\tilde{s}_t = \{0.15, 0.20, 0.25, 0.35, 0.50, 0.60\}$ for 2019–2024. The index includes both hard institutional events (legislation, ISO publications, formal methodology adoption) and soft bilateral signals (memoranda of understanding). Table 5 reports robustness to excluding soft signals from the index.

³⁸Formally, these are shifts of the form $u(e) \mapsto u(e) - \Delta_u$ that move the variance environment uniformly in a neighborhood of the threshold, consistent with the framework in Section 5.3.

³⁹The option-value weight on eligibility variance is $\beta_2 = 0.45$, reflecting evidence that investors cite certification and recognition uncertainty as a primary concern. The interaction slope $\xi_{\Omega,1} = 0.35$ governs how eligibility variance modulates the effect of flexibility on the option value; together with $\xi_{\Omega,2} = 0.08$ this parameter determines the critical variance $u^*(\phi)$ at which the regime switches (Corollary 7). The weight on technical variance is held fixed. The ordering of concerns is consistent with IEA and OECD qualitative evidence on hydrogen project bankability.

Table 1: Calibration Inputs and Implied Thresholds: Australia–Japan, 2019–2024

	2019	2020	2021	2022	2023	2024
<i>Panel A: Flexibility Index Components</i>						
Spot share (ς_t)	0.22	0.21	0.20	0.19	0.18	0.18
Supplier diversification (d_t)	0.78	0.79	0.79	0.80	0.80	0.80
Regas flexibility (r_t)	0.65	0.66	0.68	0.69	0.70	0.70
Composite index (ϕ_t)	0.55	0.55	0.56	0.56	0.56	0.56
<i>Panel B: Eligibility Clarity and Variance</i>						
Clarity index ($\tilde{s}_t$)	0.15	0.20	0.25	0.35	0.50	0.60
Eligibility variance (u_t)	0.68	0.64	0.60	0.52	0.40	0.32
<i>Panel C: Implied Threshold and Regime</i>						
Threshold (e_g^*)	0.48	0.46	0.44	0.40	0.33	0.29
$\partial e_g^*/\partial \phi$	+0.01	≈ 0	≈ 0	−0.02	−0.05	−0.06
Regime	Dest.	Dest.	Bridge	Bridge	Bridge	Bridge

Notes: Panel A reports the three components of the flexibility index from GIIGNL annual reports. Panel B reports the clarity index $\tilde{s}_t$ based on cumulative policy events and the implied eligibility variance $u_t = \bar{u}(1 - \tilde{s}_t)$ with $\bar{u} = 0.80$. Panel C reports the calibrated alignment threshold $e_g^*(\phi_t, u_t)$, the marginal effect of flexibility on the threshold (computed via the quotient rule under the interaction form), and the regime classification based on $u^*(\phi_t)$ from Corollary 7.

additive separability of the option value in technical and eligibility variance ($\Omega_{u\sigma^2} \approx 0$), corresponding to case (i) of Proposition 7.⁴⁰

6.2 Results

The calibrated threshold falls from 0.48 in 2019 to 0.29 in 2024 (Table 1, Panel C). The composite flexibility index ϕ_t is nearly constant over the sample while u_t falls by more than half. This decline in eligibility variance is the largest single movement in the calibrated inputs, and the threshold decline tracks it closely. The Australia–Japan calibration therefore isolates the eligibility-clarity channel: the time-series movement in the threshold comes almost entirely from changes in u_t rather than from changes in outside-option flexibility.

The corridor crosses the regime boundary in the early 2020s. Under the interaction-form specification, $u^*(\phi_t) \approx 0.615$ at $\phi_t \approx 0.56$.⁴¹ The constructed series lies above $u^*(\phi_t)$ in 2019–2020 and below it from 2021 onward. The marginal effect of flexibility on the threshold changes sign correspondingly, from $\partial e_g^*/\partial \phi = +0.01$ in 2019 to -0.06 in 2024. Under 2019 conditions, when eligibility variance was high, greater LNG flexibility delayed commitment by preserving the outside option; under 2024 conditions, with eligibility variance substantially reduced, the same flexibility supports commitment by providing

⁴⁰Calibrating the cross-partial $\Omega_{u\sigma^2}$ would require data on how the value of waiting responds to joint changes in technical and eligibility uncertainty, data that are not currently available for hydrogen projects. If the option value exhibits complementarity ($\Omega_{u\sigma^2} > 0$), bundled interventions would show diminishing returns, reinforcing the case for prioritizing eligibility variance when it is the dominant source of delay.

⁴¹See Theorem 5 and Corollary 7. Under the interaction form, $D(\phi) = D_0 + \xi_{\Omega,1}\eta\phi$ and the cutoff $u^*(\phi)$ is computed from the quotient-rule formula (46) with the calibrated baseline parameters. The numerical value uses $\xi_{\Omega,2} = 0.08$.




Figure 1: **Marginal Effect of Flexibility on Alignment Threshold: Australia–Japan (2019–2024).** The heatmap shows $\partial e_g^*(\phi, u)/\partial \phi$ as a function of incumbent flexibility (ϕ , horizontal axis) and eligibility variance (u , vertical axis). Positive values (warm colors) indicate the Destination regime; negative values (cool colors) indicate the Bridge regime. The dashed curve plots the regime boundary $u^*(\phi)$. Circles mark the Australia–Japan positions from 2019 (top) to 2024 (bottom). The near-vertical trajectory reflects that flexibility changed little while eligibility variance fell substantially.

insurance against downside risk.

The Australia–Japan trajectory in (ϕ, u) space is nearly vertical (Figure 1): flexibility changed little while eligibility variance fell sharply. The trajectory traces a path that is almost entirely a variance story rather than a flexibility story.

The continued absence of Australia–Japan hydrogen FIDs is consistent with these calibrated thresholds. Although the threshold has fallen from 0.48 to 0.29, the calibrated parameters imply a self-enforcing delay outcome at current clarity levels: the exporter does not find it privately profitable to supply the required alignment because the surplus from the relationship falls short of the convex alignment cost, so equilibrium effort remains at zero. Proposition 1 shows that delay under missing transfers can coexist with positive bilateral gains from early commitment under a suitable transfer.

6.3 Policy scenarios

Mutual recognition delivers the largest threshold reduction among the calibrated instruments ($\Delta e_g^* = -0.06$), reducing the threshold from 0.29 to 0.23 by guaranteeing that Australia’s hydrogen certification will be accepted under Japan’s accounting rules (Table 2). Full harmonization (clarity = 0.80) produces a larger effect (-0.08) but is best

Table 2: Policy Scenarios: Effects on Alignment Threshold Relative to 2024 Baseline

Scenario	Clarity	u	ϕ	e_g^*	Δe_g^*
Baseline (2024)	0.60	0.32	0.56	0.29	—
<i>Eligibility clarity instruments</i>					
MRA (+0.15 clarity)	0.75	0.20	0.56	0.23	−0.06
Full harmonization	0.80	0.16	0.56	0.21	−0.08
<i>Incumbent flexibility</i>					
Reduced LNG flexibility	0.60	0.32	0.46	0.30	+0.01
<i>Combined instruments</i>					
MRA + reduced flexibility	0.75	0.20	0.46	0.24	−0.05

Notes: Effects are computed under the additive specification ($\Omega_{u\sigma^2} \approx 0$); under complementarity, combined scenarios would show smaller gains than the sum of stand-alone effects (Proposition 7). Eligibility variance is $u = 0.80 \times (1 - s)$. The baseline calibration quantifies instruments that operate through eligibility variance or incumbent flexibility.

read as a limiting benchmark. In the bridge regime, reducing LNG flexibility slightly raises the threshold (+0.01) by reducing the insurance value of incumbent supply. Combining MRA with flexibility reduction yields −0.05, similar to MRA alone.

Reductions in eligibility variance lower the threshold more than flexibility restrictions under the 2024 bridge-regime baseline. MRA operates through the option-value channel identified in Proposition 3: it compresses eligibility variance, which directly reduces the value of waiting. This channel is distinct from cost-competitiveness arguments, which affect a different component of the commitment condition. The threshold-reducing effect of clarity reform is amplified by incumbent flexibility, consistent with Corollary 6. Under the LQ interaction form, the marginal clarity sensitivity $|\partial e_g^*/\partial s| = (\beta_2 + \xi_{\Omega,1}\phi)\bar{u}/D(\phi)$ is 4.4 percent larger at the baseline flexibility ($\phi = 0.56$, $|\partial e_g^*/\partial s| = 0.413$) than at the reduced-flexibility counterfactual ($\phi = 0.46$, $|\partial e_g^*/\partial s| = 0.396$).⁴² The amplification is structurally present but quantitatively modest in this corridor because incumbent flexibility varied little over the sample period. In corridors with greater flexibility variation, the interaction identified in Corollary 6 would have larger quantitative effect.

Cost competitiveness accelerates commitment only when it represents a present advantage.⁴³ When the exporter’s positioning emphasizes future cost declines and abundant supply rather than a present cost wedge, it raises the value of waiting and compresses first-mover rents, both of which raise the threshold (Corollary 4). Cost-competitive messaging and continued delay are therefore consistent when the certification uncertainty that governs the importer’s commitment margin remains unresolved.

MRA reduces the threshold from 0.29 to 0.23. Whether the exporter can privately supply that alignment depends on the implementation margin, which the threshold analysis alone does not settle.

⁴²The numerator $\beta_2 + \xi_{\Omega,1}\phi$ accounts for a 5.7 percent difference (0.646 versus 0.611); the denominator $D(\phi)$, which also increases with ϕ , attenuates this to 4.4 percent in the net sensitivity. See Appendix F.

⁴³In the model, a current cost advantage raises the static gain A through $F_I^g(e)$. Expected future cost declines correspond to the technological optimism parameter ξ in Corollary 4 and raise the option value of waiting Ω . Perceptions that capacity is not scarce reduce foregone first-mover rents δ_I^g . All three objects enter the threshold through distinct channels.

6.4 Threshold and implementation margins

The commitment threshold e_g^* is the minimum alignment at which the importer commits. Whether the exporter can privately supply that alignment is a separate question, governed by the inducement condition $B_E \geq C_E(e_g^*)$, where B_E is the exporter's net surplus and $C_E(e) = \frac{1}{2}ke^2$ is the convex alignment cost (Proposition 1). Certification clarity lowers the threshold; cost-sharing instruments relax the inducement condition. Evaluating both requires separating the two margins.

The policy scenarios in Table 2 quantify how certification clarity and incumbent flexibility affect the threshold. Carbon-price support instruments, including contracts for difference and price-gap subsidies, are absent from that table not because they are irrelevant to the threshold margin, but because the baseline calibration does not pin down the foregone-rent channel needed to quantify their threshold effect (Section 6.5 reports an equivalent carbon-price comparison under illustrative foregone-rent sensitivities). A subsidy share ω relaxes the inducement condition to $B_E \geq (1 - \omega)C_E(e_g^*)$ without changing the threshold itself. Define the surplus-cost ratio $\chi \equiv 2B_E/k$.

The implementability condition becomes

$$\chi \geq (1 - \omega) e_g^{*2}, \quad (20)$$

and the minimum subsidy share required for a relationship with surplus-cost ratio χ to implement the threshold is

$$\omega^{\text{req}}(\chi, e_g^*) = \max \left\{ 0, 1 - \frac{\chi}{e_g^{*2}} \right\}. \quad (21)$$

Certification clarity and cost sharing lower the required subsidy through different channels. Certification clarity lowers e_g^* , which reduces e_g^{*2} and therefore lowers ω^{req} for any given χ . Cost-sharing instruments provide ω directly, relaxing the constraint without moving e_g^* . Table 3 reports ω^{req} for selected surplus-cost benchmarks under baseline and MRA scenarios. Mutual recognition weakly reduces the required subsidy share at every χ ; the reduction is strict when χ lies between the post-MRA and pre-MRA implementation cutoffs. A project with $\chi = 0.06$, for example, requires a 29.7 percent subsidy share at baseline but becomes privately implementable after MRA.

MRA raises the fraction of privately implementable projects at every level of cost sharing. Suppose the surplus-cost ratio varies across potential projects, with χ drawn from a uniform distribution on $[0, \bar{\chi}]$. At any given subsidy share ω , the fraction of projects that can implement the threshold is $I(\omega \mid e_g^*) = \Pr(\chi \geq (1 - \omega)e_g^{*2})$. Figure 2 plots this implementability rate as a function of ω for each policy scenario. Certification clarity shifts the curve upward by reducing the implementation cutoff; cost-sharing instruments move along the curve by raising ω . The two instruments operate on different margins and are complements: MRA expands the set of privately implementable projects at every ω , while cost sharing brings in projects whose surplus falls short of the threshold cost.

The level of implementability depends on the support of the χ distribution, which is not calibrated. Robustness to the support choice confirms that the structural relationship is stable. MRA raises zero-subsidy implementability and weakly reduces the required subsidy share across all support specifications tested ($\bar{\chi} \in \{0.10, 0.15, 0.20, 0.25\}$). The reduction in required subsidy is strict for tighter supports and becomes zero once the

Table 3: Required Subsidy Share by Surplus-Cost Ratio and Policy Scenario

χ	Baseline ($e_g^* = 0.29$)		After MRA ($e_g^* = 0.23$)	
	ω^{req}	Implementable?	ω^{req}	Implementable?
0.02	76.6%	No	62.2%	No
0.04	53.1%	No	24.5%	No
0.06	29.7%	No	0%	Yes
0.08	6.3%	No	0%	Yes
0.10	0%	Yes	0%	Yes

Notes: $\chi = 2B_E/k$ is the exporter’s surplus-cost ratio in squared-effort units. “Implementable” means $\omega^{\text{req}} = 0$: the exporter can privately supply the threshold without cost sharing. MRA lowers e_g^* from 0.29 to 0.23 (Table 2), reducing e_g^{*2} from 0.085 to 0.053 and weakly lowering ω^{req} at every χ . A project with $\chi = 0.06$ requires a 29.7% subsidy share at baseline but becomes privately implementable after MRA.

median surplus-cost ratio already exceeds both implementation cutoffs.⁴⁴

6.5 Equivalent carbon-price comparison

The policy scenarios do not quantify carbon pricing on the threshold margin because the foregone-rent channel is not calibrated. A useful benchmark is the carbon-price increase that would generate the same threshold reduction as MRA. Under the interaction-form LQ specification,

$$\Delta \bar{\tau}^{\text{eq}} = \frac{(\beta_2 + \xi_{\Omega,1}\phi) \Delta u^{\text{MRA}}}{|\iota^g - \theta_\tau(\bar{\tau})|}, \quad (22)$$

where $\theta_\tau(\bar{\tau}) \equiv \delta_{I,0,\bar{\tau}}^g(\bar{\tau})$ is the foregone-rent sensitivity of the outside option to the carbon price level (Appendix H). The numerator is the product of the threshold’s variance sensitivity and the MRA-induced variance reduction. Under the maintained calibration mapping, $\Delta \bar{\tau}^{\text{eq}}$ is invariant to the units in which eligibility variance is measured.⁴⁵

The baseline calibration does not pin down θ_τ . Table 4 reports $\Delta \bar{\tau}^{\text{eq}}$ for illustrative values. When the foregone-rent channel is weak, matching the MRA effect requires a large carbon-price increase; when it is strong, a small increase suffices. The decisive empirical unknown on the carbon-pricing side is the foregone-rent sensitivity of the incumbent outside option.

6.6 Robustness

The regime classification and policy rankings depend on structural parameters chosen for internal consistency rather than estimated. I draw the key structural parameters independently from uniform distributions centered on their baseline values with width ± 50 percent and recompute the regime classification and policy effects.⁴⁶

⁴⁴At $\bar{\chi} = 0.10$, zero-subsidy implementability rises from 14.7 percent (baseline) to 47.0 percent (MRA); at $\bar{\chi} = 0.25$, from 65.9 percent to 78.8 percent. The direction never reverses. The implementability rate should not be read as a prediction of the FID rate. Reaching FID requires conditions beyond implementability, including demand materialization, board-level timing, and financing.

⁴⁵Under $u' = cu$, the variance-sensitivity coefficient scales as $(\beta_2 + \xi_{\Omega,1}\phi)/c$ while the MRA-induced variance shift scales as $c \Delta u^{\text{MRA}}$. Their product is invariant.

⁴⁶The parameters varied are $\xi_{\Omega,0}$, $\xi_{\Omega,1}$, $\rho'_0(\phi)$, ξ_δ , β_2 , η , and $\xi_{\Omega,2}$. Draws that violate maintained sign restrictions (e.g., $\xi_{\Omega,1} \leq 0$) or produce thresholds outside $[0, 1]$ are discarded and redrawn. I retain 1,000




Figure 2: Implementability Rate by Policy Scenario. The fraction of projects that can implement the alignment threshold as a function of the cost-sharing rate ω , under $\chi \sim \text{Uniform}[0, 0.15]$. Certification clarity instruments (MRA, full harmonization) shift curves upward by reducing e_g^{*2} . Cost-sharing instruments (CfD, co-financing) move along the horizontal axis by providing ω . The two margins are complements. MRA raises implementability at every ω , while CfD expands the implementable set for a given threshold.

The regime transition pattern is robust to joint variation in seven structural parameters by ± 50 percent. The probability of Bridge classification rises from 41 percent in 2019 to 89 percent in 2024 and crossed 50 percent in 2021.⁴⁷

Scaling the annual clarity increments by ± 30 percent generates low, central, and high clarity paths. Under the central path, u_t falls below $u^*(\phi_{2024}) \approx 0.615$ in 2021; under the high path also in 2021; under the low path in 2022. The qualitative pattern is preserved across event-coding variation.

The regime classification and policy rankings are preserved under all event-construction variations tested (Table 5).⁴⁸ The table reports three exercises: dropping individual annual event blocks one at a time, replacing judgmental weights with equal increments, and excluding the soft bilateral signal. The hardest case, dropping ISO 19870 (+0.15), raises the 2024 threshold from 0.29 to 0.35 without altering the crossing year or policy rankings. The largest effect of any single perturbation is a one-year shift in the crossing date.

The policy scenario rankings are fully robust. Across all 1,000 draws, the MRA effect exceeds the flexibility effect, with a median of -0.062 and an interquartile range

valid draws.

⁴⁷The conditional transition-year IQR among draws that exhibit a within-sample switch is 2022–2023. Approximately 41 percent of draws are already in Bridge at 2019 and approximately 11 percent remain in Destination through 2024.

⁴⁸Because the calibration is annual, the two 2024 milestones (JP GX framework and AUS-JPN hydrogen MoU) enter as a single 2024 event block. Splitting them would not affect any within-sample transition before 2024 and would only weaken the leave-one-out perturbation.

Table 4: MRA-Equivalent Carbon Price Increase as a Share of the Baseline Carbon Price

	Weak	Moderate	Strong	Very strong
Foregone-rent sensitivity (θ_τ)	0.02	0.05	0.10	0.20
MRA-equivalent $\Delta\bar{\tau}^{\text{eq}}/\bar{\tau}$	26%	6%	3%	1%

Notes: Each column represents a scenario for the foregone-rent channel strength (Appendix H). The sensitivity $\theta_\tau \equiv \delta_{T,0,\bar{\tau}}^g$ is not calibrated in the baseline; values span a range from weak to strong carbon-price timing levers. The MRA-equivalent carbon price increase is computed from (22) at the 2024 baseline ($\phi = 0.56$, $\beta_2 = 0.45$, $\xi_{\Omega,1} = 0.35$, $\iota^g = 0.01$, $\Delta u = 0.12$). The denominator uses $|\iota^g - \theta_\tau|$ rather than θ_τ alone; the distinction matters when θ_τ is close to ι^g .

of $[-0.071, -0.053]$.⁴⁹ The maintained sign restrictions $\Omega_\phi \geq 0$ and $\Omega_{\phi u} > 0$ hold over the plotted range, and the variance-dominance condition holds in all valid draws, with both the numerator $\beta_2 + \xi_{\Omega,1}\phi$ and the denominator $D(\phi)$ remaining positive over the admissible parameter draws. The calibration relies on an event-based clarity proxy whose construction is necessarily judgmental; Section 7 provides an independent descriptive plausibility check using project-level data that does not depend on the clarity index or its construction. The calibration nonetheless identifies a clear ordering of policy leverage in this corridor: certification clarity first, cost sharing second, flexibility restrictions last.

7 Descriptive evidence

The theory predicts that eligibility uncertainty delays export commitment, and that the delay should be more severe for technologies with less resolved cross-border certification. These predictions can be checked against the global hydrogen project pipeline independently of the calibration and its clarity index.

The comparison between electrolysis and CCUS provides a natural test. Electrolysis-based hydrogen faces substantial unresolved questions about cross-border certification: temporal matching between renewable generation and electrolyzer operation remains contested, additional requirements differ across jurisdictions, and grid emission accounting lacks international consensus. Hydrogen produced from natural gas with carbon capture faces relatively fewer such questions, with more established lifecycle boundaries and more mature verification protocols. If certification uncertainty is the operative friction, the export penalty should be concentrated in electrolysis.

I use the IEA Hydrogen Production Projects Database (September 2025) (International Energy Agency, 2025b), which records project status, technology, capacity, and product type for hydrogen projects worldwide. The analysis focuses on the non-China pipeline: projects at Concept, Feasibility Study, or FID/Construction stages, excluding China where state-directed investment makes bilateral eligibility uncertainty less relevant.⁵⁰ Export orientation is identified through ammonia product listing, the dominant carrier for cross-border hydrogen trade. Because this proxy is imperfect, the logit specification replaces it with a hand-coded export classification for CCUS projects, for which complete coverage is available.⁵¹

⁴⁹The flexibility effect has an IQR of $[+0.003, +0.010]$.

⁵⁰The pipeline sample contains 1,858 non-China projects. For the logit analysis, I restrict to electrolysis and CCUS projects with non-missing capacity data ($N = 1,608$). Appendix I provides further detail on the sample.

⁵¹The ammonia proxy agrees with hand-coded classification for 70.8 percent of the 528 projects where

Table 5: Event-Construction Robustness

	$\tilde{s}_{2024}$	u_{2024}	e_g^*	Regime	Crossing
<i>Panel A: Leave-one-event-out</i>					
Baseline	0.60	0.32	0.29	Bridge	2021
Drop: JP Basic H2 Strategy (−0.05)	0.55	0.36	0.31	Bridge	2022
Drop: JP Green Growth (−0.05)	0.55	0.36	0.31	Bridge	2022
Drop: IPHE GHG draft (−0.10)	0.50	0.40	0.33	Bridge	2021
Drop: ISO 19870 + harmonization (−0.15)	0.45	0.44	0.35	Bridge	2021
Drop: JP GX + AUS-JPN MoU (−0.10)	0.50	0.40	0.33	Bridge	2021
<i>Panel B: Equal-weight index</i>					
Judgmental weights	0.60	0.32	0.29	Bridge	2021
Equal weights (0.09 each)	0.60	0.32	0.29	Bridge	2020
<i>Panel C: Soft-signal exclusion</i>					
Baseline	0.60	0.32	0.29	Bridge	2021
Exclude AUS-JPN MoU (−0.05)	0.55	0.36	0.31	Bridge	2021

Notes: Panel A drops each annual event block one at a time; the parenthetical shows the removed increment. Panel B replaces the judgmental increments (+0.05, +0.05, +0.10, +0.15, +0.10) with equal increments of +0.09, preserving the total improvement of 0.45. Panel C excludes the soft bilateral signal (AUS-JPN hydrogen MoU) from the 2024 block, retaining only the hard institutional event (JP GX framework). Regime classification uses $u^*(\phi_{2024}) \approx 0.615$ from Corollary 7. Crossing denotes the first year in which $u_t < u^*(\phi_t)$. In all cases, the MRA effect exceeds the flexibility effect in absolute value.

Export-oriented electrolysis projects reach FID at roughly one-quarter the domestic rate. Among the 1,666 electrolysis projects in the non-China pipeline, 12.0 percent of domestic projects have reached FID by September 2025, compared to 3.1 percent of export-oriented projects (Table 6). The nine-percentage-point gap is statistically significant and economically large. For CCUS, the ammonia-based classification shows a reversed pattern, with export projects reaching FID at higher rates than domestic. As documented below, this reversal is sensitive to how export orientation is measured and should not be given the same interpretive weight as the electrolysis result.

The technology-specific export penalty survives controls for project scale. Table 7 reports logit estimates for FID status in the pipeline, with export orientation, technology, their interaction, and log capacity as covariates. The sample comprises 1,608 electrolysis and CCUS projects with non-missing capacity, of which 170 have reached FID. For CCUS projects, the export variable uses a hand-coded classification that covers all CCUS pipeline projects; for electrolysis projects, which are too numerous for complete hand-coding, the ammonia proxy is retained.⁵² The export penalty is more severe for electroly-

both measures are available. For electrolysis, the ammonia proxy overstates the number of export projects but preserves the direction of the FID gap. For CCUS, the proxy misclassifies a substantial share of domestic ammonia-for-fertilizer projects as export-oriented, which inflates the apparent CCUS export FID rate. Section 7.1 and Appendix I document this in detail.

⁵²Appendix I reports the same specification using the ammonia proxy for all technologies. The electrolysis export effect is virtually unchanged (−0.60 versus −0.61); the CCUS export coefficient is larger under the ammonia proxy (+1.20 versus +0.83), consistent with the proxy’s misclassification of domestic ammonia producers. Appendix I also reports a specification using the hand-coded flag for both technologies on the coded subsample ($N = 459$, 20 FID events); the electrolysis export effect is −1.00 and retains the same sign, though precision is limited.

Table 6: FID Rates by Technology and Export Orientation (Pipeline Projects)

	Domestic	Export	Difference	p -value
All technologies	11.8% ($N = 1,419$)	4.3% ($N = 439$)	+7.4 pp	<0.001
Electrolysis	12.0% ($N = 1,275$)	3.1% ($N = 391$)	+8.9 pp	<0.001
CCUS	7.4% ($N = 95$)	15.6% ($N = 45$)	−8.2 pp	0.143

Notes: IEA Hydrogen Production Projects Database, September 2025. Pipeline projects (Concept, Feasibility Study, FID/Construction) excluding China. FID denotes projects that have reached FID or begun construction by September 2025. Export orientation identified by ammonia product listing. p -values from Pearson χ^2 test (all technologies, electrolysis) and Fisher’s exact test (CCUS).

Table 7: Logit: FID Status in the September 2025 Pipeline

	(1)
Export	0.83 (0.70)
Electrolysis	−0.80* (0.45)
Export $\times$ Electrolysis	−1.43* (0.77)
Log capacity	−0.49*** (0.04)
N	1,608
FID events	170

Notes: Dependent variable equals one if project status is FID/Construction in September 2025. Sample restricted to non-China pipeline projects (Concept, Feasibility Study, FID/Construction) with non-missing capacity data. Export orientation for CCUS projects uses hand-coded classification covering all CCUS pipeline projects; for electrolysis projects, export orientation is identified by ammonia product listing. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

sis than for CCUS: the interaction term, which measures the differential penalty, is -1.43 ($p = 0.064$). The electrolysis-specific export effect implied by this interaction is -0.61 in log-odds ($p = 0.074$), corresponding to a controlled FID probability of 6.5 percent for export electrolysis at representative capacity compared to 10.7 percent for domestic electrolysis.⁵³ The CCUS-specific export effect is positive but imprecisely estimated and statistically indistinguishable from zero.

⁵³Predicted probabilities are evaluated at sample-mean log capacity. CCUS-side predicted probabilities are not reported in the text because capacity data are missing for 45 of 140 CCUS pipeline projects, and the missing observations are disproportionately domestic (80 percent). This compositional difference between the rate-table sample and the logit sample affects the CCUS point estimates; Appendix I documents this in detail.

7.1 Proxy validation

The electrolysis export penalty is not an artifact of the ammonia proxy. I hand-coded export orientation for 528 projects spanning both technologies, covering all 140 CCUS pipeline projects and 364 electrolysis pipeline projects. The ammonia proxy and hand-coded classification agree for 70.8 percent of projects. Among hand-coded electrolysis projects, the FID rate for export projects is 1.6 percent, compared to 4.7 percent for domestic projects. The gap is smaller in absolute terms than in the full ammonia-proxy sample, but the direction and the relative magnitude are preserved. For CCUS, the ammonia proxy’s apparent reversal disappears under clean coding: hand-coded export CCUS reaches FID at 9.8 percent, virtually identical to the domestic rate of 10.1 percent. The sharper CCUS contrast under the ammonia proxy is best interpreted as reflecting the misclassification of domestic projects that produce ammonia for fertilizer as export-oriented, not a genuine technology-specific pattern.

These patterns are descriptive. The ideal test would exploit exogenous variation in certification clarity across export corridors; no such variation exists in the current data. Export electrolysis projects may differ from domestic projects in developer characteristics, financing horizons, or scale ambitions that independently delay FID, and the logit controls for capacity but not for these unobserved dimensions.⁵⁴ The technology-specific pattern is nevertheless difficult to reconcile with explanations that predict a uniform export penalty, because such explanations cannot account for the absence of a comparable CCUS gap under clean classification.

8 Conclusion

Certified green commodities are credence goods. Trade in such goods requires cross-jurisdictional recognition of emissions attributes, yet the recognition environment can remain unsettled for long periods. This paper develops a bilateral real-options framework in which that unsettled environment enters as eligibility uncertainty, while incumbent flexibility enters as an outside option that shapes the timing problem. The framework separates the value of early commitment from the ability to delay it, and places certification policy on the same footing as familiar price and contracting margins.

The framework yields three results. Recognition uncertainty operates as a distinct timing margin: eligibility variance raises the option value of waiting even when carbon prices and technology costs are held fixed, because eligibility determines whether future revenues can be legally and commercially realized. The effect of incumbent flexibility on commitment is not monotonic but regime-dependent: the same flexibility that delays commitment under high eligibility variance can accelerate it under low variance, reframing the bridge-versus-destination debate as an endogenous outcome of certification clarity. For certified goods with near-zero embedded emissions, reductions in eligibility variance have larger effects on the commitment threshold than mean-price instruments when the option-value channel dominates the expected-charge channel.

The Australia–Japan calibration illustrates these mechanisms. A clarity index constructed from hard institutional developments—ISO publications, Japanese legislative milestones, and bilateral recognition events—traces the evolution of eligibility variance from 2019 to 2024; robustness checks exclude soft bilateral signals. Under earlier recognition conditions, flexibility delays commitment; under current clarity levels, the bilateral

⁵⁴Appendix I discusses these alternative explanations in detail.

relationship has moved into the bridge region, and LNG flexibility can support the transition by reducing downside risk. Descriptive evidence from the global hydrogen project pipeline points in the same direction: export-oriented electrolysis projects face high eligibility uncertainty and reach FID at markedly lower rates than domestic projects or CCUS exports with more established certification protocols. The descriptive pattern is more consistent with a certification-driven timing channel than with a generic export or scale friction.

Mutual recognition, harmonized methodology, and stable border-treatment criteria act on that state variable through reinforcing channels: mutual recognition primarily compresses eligibility variance and operates through the option-value channel, while threshold alignment primarily improves recognition probability and operates through the static-payoff channel. These instruments have greater leverage in corridors where incumbent flexibility is higher, because flexibility amplifies the option value that variance reduction compresses. When the option-value channel dominates the expected-charge channel, these instruments can move the timing margin more than mean-price support. The calibration maps the MRA effect into a carbon-price-equivalent change whose magnitude depends on the foregone-rent sensitivity of the incumbent outside option, a parameter that future empirical work should quantify.

The mechanism extends beyond green commodity trade. The timing margin analyzed here—irreversible commitment under policy-contingent eligibility with a flexible outside option—arises wherever investment requires regulatory recognition that remains unsettled across jurisdictions. The incumbent fossil system is not intrinsically a bridge or a destination. It becomes a bridge when clarity is high enough that flexibility insures rather than substitutes for commitment, and it becomes a destination when clarity remains low enough that flexibility preserves delay. Cross-border pharmaceutical approval, mutual recognition of financial regulatory standards, and agricultural certification for export markets share the same primitive: investors cannot know *ex ante* whether their product will qualify under the destination’s rules, and they retain a fallback while waiting for resolution. Several dimensions of the present setting point to natural extensions: the political economy of recognition policy, transfer instruments that bundle offtake guarantees with diplomatic commitments, carbon-price risk through an explicit foregone-rent channel, and multi-importer competition with dynamic reputation effects. These would enrich the institutional mapping, but the central message is already clear. In markets for certified goods, commitment requires more than technology readiness and price support. It requires a recognition environment that makes eligibility predictable enough for private parties to treat the green attribute as tradable, bankable, and contractible.

References

- Abel, A. B., Dixit, A. K., Eberly, J. C., and Pindyck, R. S. (1996). Options, the value of capital, and investment. *The Quarterly Journal of Economics*, 111(3):753–777.
- Agency for Natural Resources and Energy, Ministry of Economy, Trade and Industry, Government of Japan (2024). Hydrogen society promotion act enacted. toward a forthcoming hydrogen-based society. part 2: Utilization of clean hydrogen. Policy explainer.
- Aldy, J. E. and Armitage, S. (2020). The cost-effectiveness implications of carbon price certainty. *AEA Papers and Proceedings*, 110:113–118.
- Antweiler, W. and Schlund, D. (2024). The emerging international trade in hydrogen: Environmental policies, innovation, and trade dynamics. *Journal of Environmental Economics and Management*, 127:103035.
- Arrow, K. J. and Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 88(2):312–319.
- Athey, S. (2002). Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223.
- Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106.
- Besley, T. and Persson, T. (2023). The political economics of green transitions. *The Quarterly Journal of Economics*, 138(3):1863–1906.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.
- Blyth, W., Bradley, R., Bunn, D., Clarke, C., Wilson, T., and Yang, M. (2007). Investment risks under uncertain climate change policy. *Energy Policy*, 35(11):5766–5773.
- Boomsma, T. K., Meade, N., and Fleten, S.-E. (2012). Renewable energy investments under different support schemes: A real options approach. *European Journal of Operational Research*, 220(1):225–237.
- Clausing, K. A. and Wolfram, C. (2023). Carbon border adjustments, climate clubs, and subsidy races when climate policies vary. *Journal of Economic Perspectives*, 37(3):137–162.
- Dalby, P. A., Gillerhaugen, G. R., Hagspiel, V., Leth-Olsen, T., and Thijssen, J. J. (2018). Green investment under policy uncertainty and bayesian learning. *Energy*, 161:1262–1281.
- Darby, M. R. and Karni, E. (1973). Free competition and the optimal amount of fraud. *Journal of Law and Economics*, 16(1):67–88.
- Delfgaauw, J. and Swank, O. (2025). The political economy of (lacking) commitment to green policies. *Journal of Environmental Economics and Management*, 130:103133.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

- Dulleck, U. and Kerschbamer, R. (2006). On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic Literature*, 44(1):5–42.
- Dulleck, U., Kerschbamer, R., and Sutter, M. (2011). The economics of credence goods: An experiment on the role of liability, verifiability, reputation, and competition. *American Economic Review*, 101(2):526–555.
- European Commission (2023). Carbon border adjustment mechanism. Official webpage.
- European Commission (2024). Hydrogen value chain – approved IPCEIs. Competition Policy webpage. Accessed March 2026.
- Fischer, C. and Lyon, T. P. (2014). Competing environmental labels. *Journal of Economics and Management Strategy*, 23(3):692–716.
- Flora, M. and Tankov, P. (2023). Green investment and asset stranding under transition scenario uncertainty. *Energy Economics*, 124:106773.
- Fuchs, M., Stroebel, J., and Terstegge, J. (2024). Carbon VIX: Carbon price uncertainty and decarbonization investments. NBER Working Paper 32937.
- Gale, F., Goodwin, D., Lovell, H., Murphy-Gregory, H., Beasy, K., and Schoen, M. (2024). Renewable hydrogen standards, certifications, and labels: A state-of-the-art review from a sustainability systems governance perspective. *International Journal of Hydrogen Energy*, 59:654–667.
- Ganslandt, M. and Markusen, J. R. (2001). Standards and related regulations in international trade: A modeling approach. *Journal of International Economics*, 54(2):429–458.
- GIIGNL (2024). The LNG industry GIIGNL annual report 2024. Technical report, International Group of Liquefied Natural Gas Importers.
- GIIGNL (2025). The LNG industry GIIGNL annual report 2025. Technical report, International Group of Liquefied Natural Gas Importers.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.
- Handley, K. and Limaõ, N. (2015). Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy*, 7(4):189–222.
- Handley, K. and Limaõ, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review*, 107(9):2731–2783.
- Harbaugh, R., Maxwell, J. W., and Roussillon, B. (2011). Label confusion: The Groucho effect of uncertain standards. *Management Science*, 57(9):1512–1527.
- Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.
- Hartley, P. R. (2015). The future of long-term LNG contracts. *The Energy Journal*, 36(3):209–233.

- Hassett, K. A. and Metcalf, G. E. (1999). Investment with uncertain tax policy: Does random tax policy discourage investment? *The Economic Journal*, 109(457):372–393.
- Henry, C. (1974). Option values in the economics of irreplaceable assets. *The Review of Economic Studies*, 41(5):89–104.
- Howarth, R. W. (2014). A bridge to nowhere: methane emissions and the greenhouse gas footprint of natural gas. *Energy Science & Engineering*, 2(2):47–60.
- Hussain, M. and Syron, E. (2026). Global renewable hydrogen certification: Frameworks, gaps, and policies. *Energy Policy*, 213:115212.
- Hydrogen Council and McKinsey & Company (2025). Global hydrogen compass 2025. Technical report, Hydrogen Council.
- International Energy Agency (2019). The role of gas in today’s energy transitions. Technical report, IEA, Paris.
- International Energy Agency (2024). Global hydrogen review 2024. Technical report, IEA, Paris.
- International Energy Agency (2025a). Global hydrogen review 2025. Technical report, IEA, Paris.
- International Energy Agency (2025b). Hydrogen production and infrastructure projects database. September 2025 release. Licence: CC BY 4.0.
- International Organization for Standardization (2023). ISO/TS 19870:2023 hydrogen technologies — methodology for determining the greenhouse gas emissions associated with the production, conditioning and transport of hydrogen to consumption gate. Technical Specification.
- International Organization for Standardization (2024). Hydrogen technologies: An overview of ISO/TS 19870. ISO Brochure. Accessed March 2026.
- International Organization for Standardization (2025). ISO 14687:2025 hydrogen fuel quality — product specification. International Standard.
- International Partnership for Hydrogen and Fuel Cells in the Economy (2023). Methodology for determining the greenhouse gas emissions associated with the production of hydrogen. Technical report, IPHE.
- International Renewable Energy Agency (2024). Renewable power generation costs in 2023. Technical report, IRENA, Abu Dhabi.
- IPCC (2022). Climate change 2022: Mitigation of climate change. contribution of working group III to the sixth assessment report of the intergovernmental panel on climate change. Technical report, Cambridge University Press, Cambridge, UK and New York, NY, USA.
- Janzwood, A. and Millar, H. (2022). Bridge fuel feuds: The competing interpretive politics of natural gas in canada. *Energy Research & Social Science*, 88:102526.

- Japan Fair Trade Commission (2017). Survey on LNG trades (summary). Technical report, Japan Fair Trade Commission.
- Japan Hydrogen Association (2024). Survey on evaluation and certification systems for clean hydrogen (final report). Technical report, Agency for Natural Resources and Energy, Ministry of Economy, Trade and Industry. In Japanese. *Reiwa 5 nendo energy jukyū kōzō kōdoka taisaku chōsa tō jigyō: Clean suiso no hyōka ninshō taisei no kentō ni kansuru chōsa*.
- Joskow, P. L. (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. *The American Economic Review*, 77(1):168–185.
- Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from Texas oil drilling. *American Economic Review*, 104(6):1698–1734.
- Kemfert, C., Präger, F., Braunger, I., Hoffart, F. M., and Brauers, H. (2022). The expansion of natural gas infrastructure puts energy transitions at risk. *Nature Energy*, 7:582–587.
- Lee, M. and Saygin, D. (2023). Financing cost impacts on cost competitiveness of green hydrogen in emerging and developing economies. Technical Report OECD Environment Working Papers, No. 227, OECD Publishing, Paris.
- Lizzeri, A. (1999). Information revelation and certification intermediaries. *The RAND Journal of Economics*, 30(2):214–231.
- McDermott, C. L., Noah, E., and Cashore, B. (2008). Differences that “matter”? a framework for comparing environmental certification standards and government policies. *Journal of Environmental Policy & Planning*, 10(1):47–70.
- McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–728.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62(1):157–180.
- Ministry of the Environment, Government of Japan (2024). Signing of mutual recognition arrangement between Japan and Indonesia regarding the joint crediting mechanism (JCM). Press Release. Accessed January 2026.
- Nishide, K. and Nomi, E. K. (2009). Regime uncertainty and optimal investment timing. *Journal of Economic Dynamics and Control*, 33(10):1796–1807.
- Nishide, K. and Yagi, K. (2016). Investment under regime uncertainty: Impact of competition and preemption. *International Journal of Industrial Organization*, 45:47–58.
- Nunn, N. (2007). Relationship-specificity, incomplete contracts, and the pattern of trade. *The Quarterly Journal of Economics*, 122(2):569–600.
- OECD (2025). What to expect from the EU carbon border adjustment mechanism? Technical Report OECD Policy Briefs, No. 15, Organisation for Economic Co-operation and Development.

- Pindyck, R. S. (1993). Investments of uncertain cost. *Journal of Financial Economics*, 34(1):53–76.
- Richstein, J. C. and Neuhoff, K. (2022). Carbon contracts-for-difference: How to de-risk innovative investments for a low-carbon industry? *iScience*, 25(8):104700.
- Rogerson, W. P. (1992). Contractual solutions to the hold-up problem. *Review of Economic Studies*, 59(4):777–794.
- Sharples, J. (2023). A brave new world? lng contracts in the context of market turbulence and an uncertain future. Technical Report NG 187, Oxford Institute for Energy Studies.
- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- White, L. V., Fazeli, R., Beck, F. J., Baldwin, K. G. H., and Li, C. (2025). Implications for cost-competitiveness of misalignment in hydrogen certification: A case study of exports from Australia to the EU. *Energy Policy*, 204:114661.
- White, L. V., Fazeli, R., Cheng, W., Aisbett, E., Beck, F. J., Baldwin, K. G., Howarth, P., and O’Neill, L. (2021). Towards emissions certification systems for international trade in hydrogen: The policy challenge of defining boundaries for emissions accounting. *Energy*, 215:119139.
- Zahur, N. B. (2025). Long-term contracts and efficiency in the liquefied natural gas industry. SSRN Working Paper.

A Proofs of main-text results

This appendix collects formal proofs of results stated in Sections 4–5.

Proof of Theorem 1. The function $\Pi(\cdot; \phi, \bar{\tau}, \kappa_E, s)$ is continuous on $[0, 1]$, strictly increasing by Assumption 7, negative at $e = 0$, and positive at $e = 1$ by Assumption 8. The intermediate value theorem implies existence of a root; strict monotonicity implies uniqueness. $\square$

Proof of Theorem 2. Fix $(\phi, \bar{\tau}, \kappa_E, s)$ and define $F(e, x) \equiv \Pi(e; \phi, \bar{\tau}, \kappa_E, s)$. The threshold satisfies $F(e_g^*(x), x) = 0$. Differentiating with respect to x and applying the implicit function theorem yields (5). $\square$

Proof of Corollary 3. For (i): Under the reduced-form baseline, u enters Π only through Ω , so $\Pi_u = -\Omega_u < 0$ by Assumption 6(i) and $\partial e_g^*/\partial u = \Omega_u/\Pi_e > 0$. When recognition probability p depends on u through the standardized distance, A acquires an indirect u -channel and the derivative becomes $(\Omega_u - A_u)/\Pi_e$.

For (ii): Higher s raises $d = -m/\sqrt{u}$ by Assumption 4(i), which raises recognition probability p . Higher p lowers expected carbon charges in A (since $\iota^{\text{def}} > \iota^g$), raising A . By Assumption 3, $u_s < 0$, so Ω falls. Both effects raise Π : $\Pi_s = A_s - \Omega_u u_s > 0$ since $A_s > 0$ and $-\Omega_u u_s > 0$. By Theorem 2, $\partial e_g^*/\partial s < 0$.

For (iii): Higher κ_E lowers d by Assumption 4(ii), which lowers p and raises expected carbon charges, so $A_{\kappa_E} < 0$. By Assumption 3, $u_{\kappa_E} > 0$, so Ω rises. Both effects lower Π : $\Pi_{\kappa_E} = A_{\kappa_E} - \Omega_u u_{\kappa_E} < 0$ since $A_{\kappa_E} < 0$ and $-\Omega_u u_{\kappa_E} < 0$. By Theorem 2, $\partial e_g^*/\partial \kappa_E > 0$. $\square$

Proof of Corollary 4. Under the stated conditions, $\Pi_\xi = A_\xi - \Omega_\xi$. The term A_ξ includes $\delta_{I,\xi}^g < 0$, so $A_\xi < 0$. Combined with $\Omega_\xi > 0$, we have $\Pi_\xi < 0$, and Theorem 2 implies $\partial e_g^*/\partial \xi > 0$. $\square$

Proof of Proposition 2. From (4), $A_{\bar{\tau}} = \Delta_{\bar{\tau}}^F - [p \iota^g + (1-p) \iota^{\text{def}}]$. Since $\Pi = A - \Omega$ and Ω does not depend directly on $\bar{\tau}$, we have $\Pi_{\bar{\tau}} = A_{\bar{\tau}}$. Applying Theorem 2 and rearranging yields (8). $\square$

Proof of Proposition 3. From (9), $\Pi_u = A_u - \Omega_u$. By Theorem 2, $\partial e_g^*/\partial u = -\Pi_u/\Pi_e = (\Omega_u - A_u)/\Pi_e$. Since $\Pi_e > 0$, the sign equals $\text{sgn}(\Omega_u - A_u)$. $\square$

Proof of Proposition 4. Equation (12) follows from dividing (11) by the absolute value of (8). For the bound, note that

$$|p \iota^g + (1-p) \iota^{\text{def}} - \Delta_{\bar{\tau}}^F| \leq |(1-p) \iota^{\text{def}} - \Delta_{\bar{\tau}}^F| + p \bar{\iota}$$

by the triangle inequality. Requiring the ratio in (12) to exceed R and substituting $\iota^g \leq \bar{\iota}$ yields (13). $\square$

Proof of Proposition 5. From $\Pi = A - \Omega(\phi, u)$ and (4), the total derivative with respect to s is

$$\Pi_s = A_s - \Omega_u \cdot u_s,$$

where $A_s = \tilde{A}_s + \bar{\tau}(\iota^{\text{def}} - \iota^g)p_s$. Since $\tilde{A}_s \geq 0$, $p_s > 0$, $\iota^{\text{def}} > \iota^g$, and $\bar{\tau} \geq 0$ (carbon prices are non-negative), we have $A_s \geq 0$. By Assumption 12(i), $u_s < 0$, and by Assumption 6(i), $\Omega_u > 0$. Therefore $-\Omega_u \cdot u_s > 0$. Since the variance channel is strictly positive and the static channel is non-negative, $\Pi_s > 0$. By Theorem 2, $\partial e_g^*/\partial s = -\Pi_s/\Pi_e < 0$. $\square$

Proof of Corollary 6. The cross-partial of the net gain with respect to ϕ and s is $\Pi_{\phi s} = A_{\phi s} - \Omega_{\phi u} \cdot u_s$. By Assumption 9, $A_{\phi s} = 0$. By Assumption 6(iii), $\Omega_{\phi u} > 0$, and by Assumption 12(i), $u_s < 0$. Therefore $\Pi_{\phi s} = -\Omega_{\phi u} \cdot u_s > 0$ for all (e, ϕ, s) . Differentiating (16) with respect to ϕ gives $\partial \Delta \Pi / \partial \phi = \int_{s_0}^{s_1} \Pi_{\phi s}(e; \phi, s) ds > 0$, since the integrand is strictly positive throughout. $\square$

B Timing robustness

The baseline model uses Stackelberg timing in which the exporter chooses alignment effort, the importer observes it, and then decides whether to commit. This appendix shows that the threshold structure is robust to alternative timing assumptions.

B.1 Within-period commitment after observing effort

Suppose the exporter and importer act within the same period, but the importer observes the exporter's effort choice before making its commitment decision. The importer's best response to any observed effort level e is to commit if and only if $\Pi(e; \phi) \geq 0$, which by single-crossing (Assumption 7) is equivalent to $e \geq e_g^*(\phi)$. The exporter, anticipating this best response, faces the same inducement problem as in the sequential game: it can either supply zero effort and accept that the importer waits, or supply the minimal inducement $e = e_g^*(\phi)$ and secure commitment. The exporter chooses the latter if and only if $\Delta_E \geq 0$. The equilibrium outcome therefore replicates the Stackelberg equilibrium, and the threshold $e_g^*(\phi)$ retains its role as the critical object.

B.2 Noisy public signals

Suppose the importer does not observe the exporter's effort e directly but instead observes a public signal $\tilde{e}$, a noisy proxy for e . Assume the signal distribution satisfies the monotone likelihood ratio property: for any $e' > e$, the likelihood ratio $f(\tilde{e} | e') / f(\tilde{e} | e)$ is strictly increasing in $\tilde{e}$.

Under the monotone likelihood ratio property, the posterior over e is increasing in $\tilde{e}$, and the importer's optimal decision takes a cutoff form in the signal. There exists a unique lower threshold $\bar{e}(\phi)$ such that the importer commits if and only if $\tilde{e} \geq \bar{e}(\phi)$. The exporter's problem becomes one of choosing effort to induce signal realizations above this cutoff at minimum cost.

When the net gain Π is affine in e (holding ϕ fixed), the cutoff in signals is equivalent to a cutoff in the posterior mean $\mu(\tilde{e}) \equiv \mathbb{E}[e | \tilde{e}]$. In that case, $\mathbb{E}[\Pi(e; \phi) | \tilde{e}] = \Pi(\mu(\tilde{e}); \phi)$, so the importer commits if and only if $\mu(\tilde{e}) \geq e_g^*(\phi)$. The comparative statics derived in the main text therefore apply to the signal-based setting with the posterior mean replacing actual effort, and all signs are preserved.

B.3 Two-sided effort

Suppose the importer can also exert due diligence effort $e_I \in [0, 1]$ that reduces uncertainties, at convex cost $C_D(e_I)$ with $C_D(0) = 0$, $C_D' > 0$, and $C_D'' > 0$. Assume each variance term, $\sigma^2(e, e_I)$ and $u(e, e_I)$, is decreasing in each argument. Assume the importer's net gain $\Pi(e, e_I; \phi)$ is decreasing in each variance term.

Importer effort e_I enters alongside exporter effort e through the variance terms in the commitment condition. Fix e_I . The single-crossing property in exporter effort is

preserved, so $\Pi(e, e_I; \phi)$ is strictly increasing in e and there exists a threshold $e_g^*(e_I, \phi)$. Fix e . The net gain is increasing in e_I , and an interior choice of e_I solves the first-order condition when the optimum lies in $(0, 1)$.

Fixing the importer's due-diligence choice, the comparative statics from the main text carry over directly. With endogenous e_I , the same directional effects obtain provided the induced change in e_I does not exactly offset the change in total uncertainty reduction implied by the primitive shift. The key structural result survives: a threshold in total uncertainty reduction exists and can be achieved through combinations of e and e_I , and that threshold retains the gap-over-slope structure characterized in the main text.

C Mean and variance channels for recognition

Proposition 6 (Channel decomposition for recognition probability). *Under Assumptions 1 and 2, the marginal effect of alignment effort on recognition probability satisfies*

$$\frac{\partial p}{\partial e} = \underbrace{-\frac{f(d)}{\sqrt{u}}}_{p_m} \cdot m_e + \underbrace{\frac{m \cdot f(d)}{2u^{3/2}}}_{p_u} \cdot u_e, \quad (23)$$

where $d \equiv -m/\sqrt{u}$. The mean channel $p_m \cdot m_e$ is unambiguously positive since $p_m < 0$ and $m_e < 0$. The variance channel $p_u \cdot u_e$ has sign equal to $-\text{sgn}(m)$:

- (i) If $m > 0$ (expected noncompliance), then $p_u > 0$ and $u_e < 0$, so variance reduction lowers recognition probability by shrinking the left tail where lucky recognition occurs.
- (ii) If $m < 0$ (expected compliance), then $p_u < 0$ and $u_e < 0$, so variance reduction raises recognition probability by shrinking the right tail where unlucky non-recognition occurs.

Proof. From (2), $p = F(d)$ where $d \equiv -m/\sqrt{u}$. Differentiating with respect to e :

$$\frac{\partial p}{\partial e} = f(d) \cdot \frac{\partial d}{\partial e} = f(d) \left[-\frac{m_e}{\sqrt{u}} + \frac{m \cdot u_e}{2u^{3/2}} \right].$$

Identifying $p_m = -f(d)/\sqrt{u} < 0$ and $p_u = m \cdot f(d)/(2u^{3/2})$, we have $\text{sgn}(p_u) = \text{sgn}(m)$. Since $u_e < 0$, the variance channel $p_u \cdot u_e$ has sign $-\text{sgn}(m)$. $\square$

D Transfers and delay inefficiency

This appendix provides the proof and interpretation of Proposition 1, stated in Section 4.2.

Proof. Fix any parameter values and $\tilde{e}$ satisfying $\Pi(\tilde{e}; \phi) > 0$ and $C_E(\tilde{e}) - B_E > 0$. If, in addition,

$$\Pi(\tilde{e}; \phi) \geq C_E(\tilde{e}) - B_E,$$

then a transfer z exists with

$$C_E(\tilde{e}) - B_E \leq z \leq \Pi(\tilde{e}; \phi).$$

Such a transfer makes the exporter's net payoff weakly higher than delay, since $B_E - C_E(\tilde{e}) + z \geq 0$, and leaves the importer weakly better off than delay, since $\Pi(\tilde{e}; \phi) - z \geq 0$. $\square$

The transfer z in Proposition 1 can be interpreted as a bilateral side payment from importer to exporter. Such a payment requires $\tilde{e} > e_g^*(\phi)$ because $\Pi(e_g^*(\phi); \phi) = 0$ leaves no importer surplus to finance a strictly positive payment while remaining strictly better off. The same formal wedge is also consistent with external co-financing (for example, government or platform funding for certification pilots or shared measurement infrastructure), which operates by lowering the exporter's effective alignment cost $C_E(\cdot)$ rather than reallocating surplus from the importer.

E Microfoundations for LNG channels

The model assumes that incumbent flexibility ϕ affects the baseline risk premium $\rho_0(\phi)$, foregone first-mover rents $\delta_I^g(\phi)$, and the option value of waiting $\Omega(\cdot; \phi)$. This appendix provides economic foundations for these reduced-form assumptions in the context of LNG trade.

E.1 Portfolio insurance and the risk premium

Let the importer discount risky payoffs at rate $r + \lambda\sigma_R$, where σ_R denotes residual revenue volatility after hedging and $\lambda > 0$ is the market price of risk. When the importer holds a portfolio of energy contracts, greater flexibility in the incumbent LNG portfolio (destination rights, volume adjustment provisions, and access to spot markets) improves hedging and reduces residual revenue volatility. Assume $\sigma_R(\phi) = \bar{\sigma}_R - \nu\phi$ for some $\nu > 0$.

The effective risk premium on the green investment is $\rho_0(\phi) = \lambda\sigma_R(\phi) = \lambda\bar{\sigma}_R - \lambda\nu\phi$. The derivative satisfies $\rho_{0,\phi} = -\lambda\nu < 0$, so greater incumbent flexibility reduces the baseline risk premium by improving portfolio insurance during the transition period.

E.2 Foregone rents and alternative hedges

Early commitment to the green supply chain confers first-mover rents: priority access to capacity, participation in emerging hubs, influence over standards, and network externalities from early coordination. Let R denote the gross first-mover rent available at $t = 1$. If the importer delays, it foregoes some of these rents, but it also retains the value of hedges available through the incumbent portfolio during the waiting period. Let $H(\phi)$ denote this hedge value, also evaluated at $t = 1$ and increasing in ϕ .

Define the net foregone rent from delaying as $\delta_I^g(\phi) \equiv R - H(\phi)$. Under the assumption that $H(\phi) = H_0 + \theta_H\phi$ with $\theta_H > 0$, define the baseline foregone rent as $\delta_{I,0}^g \equiv R - H_0$. Then $\delta_I^g(\phi) = \delta_{I,0}^g - \theta_H\phi$, which yields $\delta_{I,\phi}^g = -\theta_H < 0$: greater incumbent flexibility reduces the cost of waiting by improving hedging opportunities in the incumbent portfolio. The hedge term $H(\phi)$ captures insurance available even if the green project is not adopted, distinct from the option-value channel that operates through the distribution of $\tilde{V}$.

E.3 Option value and the choice set during waiting

The option value of waiting arises from the ability to observe how eligibility uncertainty resolves before committing.

Let $G = m(e, \kappa_E, s) + \sqrt{u(e, \kappa_E, s)}\varepsilon$ denote the recognition gap, where ε is a mean-zero, unit-variance shock with density f and cdf F . Certification clarity s , exporter consistency κ_E , and alignment effort e govern the location and dispersion objects (m, u) that summarize this eligibility uncertainty. Define $X \equiv -G$. Without loss of generality,

redefine the shock as $\varepsilon \leftarrow -\varepsilon$ (and denote its cdf and pdf again by F and f), so that $X = -m(e, \kappa_E, s) + \sqrt{u(e, \kappa_E, s)} \varepsilon$; X has location $-m$ and scale $\sqrt{u}$.

The importer commits at $t = 1$ if $X \geq K(\phi)$, where the effective threshold $K(\phi)$ satisfies $K'(\phi) < 0$: higher incumbent flexibility (spot depth, diversification, regasification slack) reduces the friction of unwinding LNG positions when transitioning to green commitment. Incumbent flexibility affects the conversion threshold $K(\phi)$ rather than the distribution of X , so ϕ operates through the threshold-crossing event rather than through the moments (m, u) . This threshold channel is distinct from the hedge $H(\phi)$, which captures deterministic continuation value absorbed into A through $\delta_I^g(\phi)$. The option value $\Omega(\phi, u)$ captures dispersion/learning value that operates only when uncertainty resolves.

Define the option value as $\Omega(\phi, u) \equiv \mathbb{E}[(X - K(\phi))_+]$. The derivative with respect to flexibility is

$$\Omega_\phi(\phi, u) = -K'(\phi) \Pr(X > K(\phi)) \geq 0.$$

Since $K'(\phi) < 0$, the inequality $\Omega_\phi \geq 0$ holds globally.

For the cross-partial, let $d \equiv (K(\phi) + m)/\sqrt{u}$. Then

$$\Omega_{\phi u}(\phi, u) = -K'(\phi) \cdot \frac{d}{2u} f(d),$$

which is positive whenever $d > 0$ and $f(d) > 0$, i.e., $K(\phi) + m > 0$, under standard location-scale families. The condition $K(\phi) + m > 0$ corresponds to the empirically relevant delay region in which expected eligibility falls short of the conversion threshold after accounting for incumbent flexibility, so waiting has positive option value. In this region, $\Omega_{\phi u} > 0$: flexibility is more valuable when variance is higher.

The second derivative $\Omega_{\phi\phi}$ depends on the functional form of $K(\phi)$ and the distribution shape; the general theory does not sign this term.

E.4 Stopping interpretation and explicit option-value properties

The preceding subsection defines the option value as a call-option expectation and derives $\Omega_\phi \geq 0$ and $\Omega_{\phi u} > 0$. This subsection adds three pieces: an explicit derivation of $\Omega_u > 0$, a 2-period stopping interpretation that gives the option-exercise structure an information-acquisition foundation, and a primitive parameter configuration under which the limit conditions in Assumption 10 are satisfied.

Explicit derivation of $\Omega_u > 0$. Recall $\Omega(\phi, u) = \mathbb{E}[(X - K(\phi))_+]$ where $X = -m + \sqrt{u} \varepsilon$ with ε having mean zero, unit variance, continuous density $f > 0$, and cdf F . Substituting $z = (x + m)/\sqrt{u}$ and writing $d \equiv (K(\phi) + m)/\sqrt{u}$ gives

$$\Omega(\phi, u) = \sqrt{u} \int_d^\infty (z - d) f(z) dz.$$

Differentiating with respect to u and using $\partial d / \partial u = -d/(2u)$ yields, after simplification,

$$\Omega_u(\phi, u) = \frac{1}{2\sqrt{u}} \int_d^\infty z f(z) dz. \quad (24)$$

In the delay region $d > 0$, the integrand is strictly positive on (d, ∞) , so $\Omega_u > 0$. More generally, for symmetric location-scale families within the class of Assumption 1,

$\mathbb{E}[z \mathbf{1}(z > d)] > 0$ for all $d \in \mathbb{R}$, so $\Omega_u > 0$ globally. This establishes Assumption 6(i) from the option-exercise primitive.

Two-period stopping interpretation. The call-option expectation $\Omega(\phi, u) = \mathbb{E}[(X - K(\phi))_+]$ arises as the continuation value in a standard 2-period stopping problem.

- *Period 0.* The importer observes the current state (ϕ, u, m) but not the realization of the eligibility state X . It chooses between committing now ($a_0 = C$) and waiting ($a_0 = W$).
- *Period 1* (reached only if waited). The eligibility state X is realized and observed. The importer commits if $X \geq K(\phi)$ and abandons otherwise.

Commitment is irreversible. During the waiting period, the importer retains the incumbent portfolio with flexibility ϕ , which provides a deterministic fallback absorbed into the static payoff A through the hedge $H(\phi)$ and foregone rents $\delta_I^g(\phi)$.

The optimal period-1 policy is the threshold rule: commit if and only if $X \geq K(\phi)$. The expected surplus from waiting, net of the deterministic fallback components already in A , is

$$\Omega(\phi, u) = \beta \mathbb{E}[(X - K(\phi))_+],$$

where $\beta \in (0, 1)$ is the discount factor (absorbed into the normalization of Ω in the main text, which sets $\beta = 1$ without loss of generality for the sign results).

The “learning” content of this environment is: at period 0, the importer faces eligibility uncertainty summarized by (m, u) ; by waiting to period 1, the actual eligibility state X is revealed. Higher u raises the dispersion of X and thus the value of learning before committing, so $\Omega_u > 0$ via (24). Higher ϕ lowers the conversion threshold $K(\phi)$ and expands the set of realizations at which commitment is worthwhile, so $\Omega_\phi \geq 0$. The interaction $\Omega_{\phi u} > 0$ arises because when u is high, the expanded acceptance region from lower $K(\phi)$ captures more probability mass in the dispersed tail.

Sufficient configuration for the crossing condition. Assumption 10 requires that both bridge and destination regimes are attainable: $A_\phi > \Omega_\phi(\phi, \underline{u})$ at low variance and $A_\phi < \Omega_\phi(\phi, \bar{u})$ at high variance. The following construction shows that such a configuration exists in the stopping environment.

Lemma 2 (Existence of crossing configuration). *In the 2-period stopping environment above, suppose:*

- (i) $0 < A_\phi < \kappa/2$ (static flexibility effect is bridge-favoring, but bounded);
- (ii) $K'(\phi) < 0$ and $|K'(\phi)| \geq \kappa > 0$;
- (iii) f is symmetric and continuous with $F(0) = 1/2$;
- (iv) $K(\phi) + m > 0$ (delay region).

Then there exist $0 < \underline{u} < \bar{u}$ such that the limit conditions in Assumption 10 hold.

Proof. From the option-exercise representation, $\Omega_\phi(\phi, u) = -K'(\phi)(1 - F(d))$ with $d = (K(\phi) + m)/\sqrt{u}$.

As $u \rightarrow 0^+$: by condition (iv), $K(\phi) + m > 0$, so $d \rightarrow +\infty$ and $1 - F(d) \rightarrow 0$. Therefore $\Omega_\phi \rightarrow 0$. Since $A_\phi > 0$ by condition (i), $A_\phi - \Omega_\phi > 0$ for sufficiently small u .

As $u \rightarrow \infty$: $d \rightarrow 0^+$, so $1 - F(d) \rightarrow 1 - F(0) = 1/2$ by symmetry (iii). Therefore $\Omega_\phi \rightarrow |K'(\phi)|/2 \geq \kappa/2$. Since $A_\phi < \kappa/2$ by condition (i), $A_\phi - \Omega_\phi < 0$ for sufficiently large u .

Since $\Omega_\phi(\phi, u)$ is continuous in u (by continuity of F and the integrand), the intermediate value theorem implies the existence of finite $0 < \underline{u} < \bar{u}$ satisfying the limit conditions in Assumption 10. $\square$

Lemma 2 is an existence result, not a general characterization. The parameter restriction $0 < A_\phi < \kappa/2$ is economically interpretable: it requires that the net static flexibility benefit (risk-premium reduction minus foregone-rent effect) is positive but smaller than half the threshold-lowering effect of flexibility on the commitment margin. The main text's Theorem 5 takes the crossing as given and derives the regime-switch structure from it; this lemma provides one primitive foundation. The 2-period structure is chosen for transparency; richer stopping environments with gradual learning would generate the same sign properties under standard regularity conditions.

F Linear–Quadratic specialization and derivations

This appendix collects the linear–quadratic specialization of the model and derives all comparative statics referenced in the main text. This specialization is used only to obtain closed forms and magnitude expressions; all sign results follow from Section 4.

F.1 LQ primitives

Impose the following functional forms:

$$\rho(e; \phi) = \rho_0(\phi) - \alpha e, \quad \text{with } \rho_0(\phi) = \bar{\rho} - \chi\phi, \quad (25)$$

$$F_I^g(e) = f_I e, \quad (26)$$

$$\delta_I^g(\phi) = \delta_{I,0}^g - \xi_\delta \phi, \quad (27)$$

$$\sigma^2(e) = \sigma_0^2 - \gamma e, \quad (28)$$

$$u(e) = u_0 - \eta e, \quad (29)$$

$$\Omega(\sigma^2, u; \phi) = \beta_1 \sigma^2 + \beta_2 u + \xi_\Omega \phi, \quad (30)$$

where all Greek parameters are positive. The parameters α , γ , and η measure how effectively exporter effort reduces risk premia, technical variance, and eligibility variance, respectively. The parameters β_1 and β_2 measure how these variances contribute to option value, and ξ_Ω measures the direct effect of incumbent flexibility on option value. The object u_0 denotes baseline eligibility variance at zero alignment effort, paralleling σ_0^2 for technical variance.

F.2 Threshold derivation

The commitment condition $\Pi(e; \phi) = 0$ can be written as

$$\rho(e; \phi) + \bar{\tau} \iota^g - F_I^g(e) - \delta_I^g(\phi) + \Omega(\sigma^2(e), u(e); \phi) = 0. \quad (31)$$

Substituting the LQ forms and collecting terms in e yields

$$(\bar{\rho} - \chi\phi - \alpha e) + \bar{\tau} \iota^g - f_I e - (\delta_{I,0}^g - \xi_\delta \phi) + \beta_1(\sigma_0^2 - \gamma e) + \beta_2(u_0 - \eta e) + \xi_\Omega \phi = 0.$$

Rearranging and solving for the threshold gives

$$e_g^*(\phi) = \frac{N(\phi)}{D}, \quad (32)$$

where the numerator, which collects the importer's reasons to wait, is

$$N(\phi) = \bar{\rho} + \bar{\tau}\iota^g + \beta_1\sigma_0^2 + \beta_2u_0 - \delta_{I,0}^g + (-\chi + \xi_\delta + \xi_\Omega)\phi, \quad (33)$$

and the denominator, which measures the closing power of alignment effort, is

$$D = \alpha + f_I + \beta_1\gamma + \beta_2\eta > 0. \quad (34)$$

F.3 Comparative statics

Under the simple specification (30), $e_g^*(\phi) = N(\phi)/D$ with D a positive constant, so the partial derivatives follow directly from differentiating $N(\phi)$. Under the interaction specification (41), the denominator $D(\phi)$ depends on ϕ ; modified comparative statics are collected in the interaction-form subsection below. The comparative statics here use increasing-differences and single-crossing logic in the sense of Topkis (1998) and Milgrom and Shannon (1994), with the uncertainty extension following Athey (2002).

The sensitivity to expected carbon price is

$$\frac{\partial e_g^*}{\partial \bar{\tau}} = \frac{\iota^g}{D}. \quad (35)$$

For certified green goods with assigned intensity $\iota^g \approx 0$, this derivative is near zero. This is a partial derivative that holds foregone rents and other primitives fixed; see Appendix H for the full analysis including the foregone-rent channel.

The sensitivity to eligibility variance is

$$\frac{\partial e_g^*}{\partial u_0} = \frac{\beta_2}{D} > 0. \quad (36)$$

Higher eligibility variance raises the threshold and delays commitment. Eligibility variance is the central carbon-related object for timing on the certification margin.

The sensitivity to technical variance is

$$\frac{\partial e_g^*}{\partial \sigma_0^2} = \frac{\beta_1}{D} > 0. \quad (37)$$

The sensitivity to incumbent flexibility under the simple LQ form (30) is

$$\frac{\partial e_g^*}{\partial \phi} = \frac{-\chi + \xi_\delta + \xi_\Omega}{D}. \quad (38)$$

The sign depends on whether risk reduction (χ) exceeds foregone rents plus option value ($\xi_\delta + \xi_\Omega$).

F.4 Expected cost drift

Let $\xi \geq 0$ index expected green cost drift, entering through foregone rents $\delta_{I,0}^g(\xi)$ and option value $\Omega(\cdot; \xi)$. From the implicit function theorem,

$$\frac{\partial e_g^*}{\partial \xi} = -\frac{\Pi_\xi}{\Pi_e}, \quad \text{where} \quad \Pi_\xi = \delta_{I,0,\xi}^g - \Omega_\xi. \quad (39)$$

Under $\delta_{I,0,\xi}^g < 0$ (first-mover rents shrink when everyone expects costs to fall) and $\Omega_\xi > 0$ (the option to wait is more valuable when future costs are expected to be lower), we have $\Pi_\xi < 0$ and hence $\partial e_g^* / \partial \xi > 0$. This establishes Corollary 4.

In an LQ specialization, suppose $\delta_{I,0}^g(\xi) = \delta_0 - \chi_\delta \xi$ and $\Omega(\cdot; \xi) = \Omega_0 + \chi_\Omega \xi$ with $\chi_\delta, \chi_\Omega > 0$. Then

$$\frac{\partial e_g^*}{\partial \xi} = \frac{\chi_\delta + \chi_\Omega}{D} > 0. \quad (40)$$

The magnitude depends on the sensitivity of foregone rents and option value to expected drift, but the sign is unambiguous under the stated conditions.

F.5 Interaction form and the bridge/destination cutoff

To allow the bridge/destination classification to depend on eligibility variance, I adopt the following parametric specification for the option-value term $\Omega(\sigma^2, u; \phi)$ in the linear-quadratic environment:

$$\Omega(\sigma^2, u; \phi) = \beta_1 \sigma^2 + \beta_2 u + (\xi_{\Omega,0} + \xi_{\Omega,1} u) \phi + \frac{1}{2} \xi_{\Omega,2} \phi^2 \quad (41)$$

where $\xi_{\Omega,0}$ is the baseline effect of flexibility on option value (which can be negative if flexibility reduces option value at low eligibility variance), $\xi_{\Omega,1} > 0$ captures increasing differences in (ϕ, u) , and $\xi_{\Omega,2} \geq 0$ captures curvature with respect to ϕ .⁵⁵ The baseline calibration uses $\xi_{\Omega,2} = 0.08$.

Under (41), the cross-term $\xi_{\Omega,1}(u_0 - \eta e)\phi$ introduces dependence of the commitment condition on the product $e\phi$. The threshold is still a ratio, but the denominator now depends on ϕ :

$$e_g^*(\phi) = \frac{N(\phi)}{D(\phi)}, \quad (42)$$

where

$$N(\phi) = \bar{\rho} + \bar{\tau} \iota^g + \beta_1 \sigma_0^2 + \beta_2 u_0 - \delta_{I,0}^g + (-\chi + \xi_\delta + \xi_{\Omega,0} + \xi_{\Omega,1} u_0) \phi + \frac{1}{2} \xi_{\Omega,2} \phi^2, \quad (43)$$

$$D(\phi) = D_0 + \xi_{\Omega,1} \eta \phi, \quad D_0 \equiv \alpha + f_I + \beta_1 \gamma + \beta_2 \eta. \quad (44)$$

The term $\xi_{\Omega,1} \eta \phi$ in the denominator captures the interaction channel: when the incumbent market is more flexible, alignment effort is more effective at reducing the threshold because effort compresses the eligibility variance that flexibility makes valuable.

⁵⁵ Although $\xi_{\Omega,0}$ is not sign-restricted, the effective marginal option value of flexibility, $\Omega_\phi = \xi_{\Omega,0} + \xi_{\Omega,1} u + \xi_{\Omega,2} \phi$, is positive over the calibrated domain ($\phi \in [0.55, 0.56]$, $u \in [0.32, 0.68]$). The maintained property $\Omega_\phi \geq 0$ from Assumption 6 therefore holds in the parameter region used for calibration.

The sensitivity to incumbent flexibility follows from the quotient rule:

$$\frac{\partial e_g^*}{\partial \phi} = \frac{N'(\phi) D(\phi) - N(\phi) D'(\phi)}{D(\phi)^2} \equiv \frac{\Lambda(\phi, u_0)}{D(\phi)^2}, \quad (45)$$

where $N'(\phi) = -\chi + \xi_\delta + \xi_{\Omega,0} + \xi_{\Omega,1}u_0 + \xi_{\Omega,2}\phi$ and $D'(\phi) = \xi_{\Omega,1}\eta$. Since $D(\phi)^2 > 0$, the sign of $\partial e_g^*/\partial \phi$ is determined by Λ .

Corollary 7 (LQ cutoff). *The numerator $\Lambda(\phi, u_0)$ is affine in u_0 for fixed ϕ . Setting $\Lambda = 0$ and solving for u_0 yields the critical eligibility variance:*

$$u^*(\phi) = \frac{\xi_{\Omega,1} \eta N_0(\phi) + (\chi - \xi_\delta - \xi_{\Omega,0} - \xi_{\Omega,2}\phi) D(\phi)}{\xi_{\Omega,1} \tilde{D}_0}, \quad (46)$$

where $N_0(\phi) \equiv N(\phi, u_0=0)$ is the u_0 -free part of the numerator and $\tilde{D}_0 \equiv D_0 - \beta_2\eta = \alpha + f_I + \beta_1\gamma$ is the portion of the denominator attributable to channels other than eligibility variance reduction. For $u_0 < u^*(\phi)$, the incumbent is in the bridge regime ($\partial e_g^*/\partial \phi < 0$); for $u_0 > u^*(\phi)$, the incumbent is in the destination regime ($\partial e_g^*/\partial \phi > 0$).

Proof. Write $N(\phi) = N_0(\phi) + (\beta_2 + \xi_{\Omega,1}\phi) u_0$ and $N'(\phi) = c_\phi + \xi_{\Omega,1} u_0$, where $c_\phi \equiv -\chi + \xi_\delta + \xi_{\Omega,0} + \xi_{\Omega,2}\phi$ collects the u_0 -free terms. Substituting into $\Lambda = N'D - ND'$ and expanding, the ϕ -dependent terms carrying u_0 cancel pairwise (the product $\xi_{\Omega,1}u_0 \cdot D(\phi)$ from the first term offsets $(\beta_2 + \xi_{\Omega,1}\phi) u_0 \cdot \xi_{\Omega,1}\eta$ from the second when the shared component $\xi_{\Omega,1}^2\eta\phi u_0$ is extracted). The remainder is $\Lambda = c_\phi D(\phi) - N_0(\phi) \xi_{\Omega,1}\eta + \xi_{\Omega,1} \tilde{D}_0 u_0$. Setting $\Lambda = 0$ and solving gives (46). $\square$

Remark 2 (Constant cutoff under $\xi_{\Omega,2} = 0$). *When $\xi_{\Omega,2} = 0$, both $N_0(\phi)$ and c_ϕ are affine in ϕ . In Λ , the product $c_\phi D(\phi)$ generates ϕ and ϕ^2 terms that cancel against $N_0(\phi) \xi_{\Omega,1}\eta$. The cutoff reduces to the constant*

$$u_0^* = \frac{N_0 \xi_{\Omega,1} \eta + (\chi - \xi_\delta - \xi_{\Omega,0}) D_0}{\xi_{\Omega,1} \tilde{D}_0}, \quad (47)$$

where $N_0 = \bar{\rho} + \bar{\tau}\iota^g + \beta_1\sigma_0^2 - \delta_{I,0}^g$, independent of ϕ . This expression exceeds the constant-denominator formula $(\chi - \xi_\delta - \xi_{\Omega,0})/\xi_{\Omega,1}$ through a scaling effect ($D_0/\tilde{D}_0 > 1$) and an additive level effect ($N_0\eta/\tilde{D}_0 > 0$), both reflecting the absorption of the eligibility-variance channel into $D(\phi)$.

Under the interaction form, the local sensitivity of the threshold to baseline eligibility variance becomes

$$\frac{\partial e_g^*}{\partial u_0} = \frac{\beta_2 + \xi_{\Omega,1}\phi}{D(\phi)}. \quad (48)$$

The numerator exceeds the simple-form value β_2 for all $\phi > 0$: incumbent flexibility amplifies the timing effect of eligibility variance through the direct complementarity channel $\xi_{\Omega,1}\phi$. Greater flexibility raises the marginal timing value of clarity-oriented policies because it increases the amount by which a reduction in baseline eligibility variance lowers the commitment threshold. This amplification does not affect the direct assigned-intensity channel alone, $\partial e_g^*/\partial \bar{\tau} = \iota^g/D(\phi)$, which remains near zero for certified goods with $\iota^g \approx 0$. Comparison with the full carbon-price timing effect requires the total derivative incorporating the foregone-rent channel (Appendix H).

F.6 Clarity–flexibility interaction

Corollary 6 establishes that clarity reform improves the net gain from commitment more at higher flexibility, for any fixed alignment effort. Under the LQ interaction form, the result extends to the equilibrium threshold. Clarity enters through $u_0 = \bar{u}(1 - s)$, so $\partial u_0 / \partial s = -\bar{u}$. From (42), the threshold sensitivity to clarity is

$$\frac{\partial e_g^*}{\partial s} = -\frac{(\beta_2 + \xi_{\Omega,1}\phi) \bar{u}}{D(\phi)} < 0. \quad (49)$$

Differentiating with respect to ϕ via the quotient rule, and noting that $D(\phi) = D_0 + \xi_{\Omega,1}\eta\phi$ does not depend on s :

$$\frac{\partial^2 e_g^*}{\partial s \partial \phi} = \frac{-\bar{u} \xi_{\Omega,1} \tilde{D}_0}{D(\phi)^2} < 0, \quad (50)$$

where $\tilde{D}_0 = \alpha + f_I + \beta_1\gamma > 0$ is the portion of the denominator attributable to channels other than eligibility variance reduction. The sign is determined by the positivity of all constituent parameters ($\bar{u} > 0$, $\xi_{\Omega,1} > 0$, $\tilde{D}_0 > 0$) and does not require additional assumptions beyond the LQ primitives.

Higher incumbent flexibility amplifies the threshold-reducing effect of clarity reform. The numerator $\xi_{\Omega,1}\tilde{D}_0$ reflects the same increasing-differences parameter that drives the regime-switch result (Corollary 7), scaled by the closing power of channels other than eligibility variance. The denominator $D(\phi)^2$ grows with ϕ , but this attenuation is dominated by the numerator effect. This confirms the general-form Corollary 6 at the equilibrium-threshold level under the LQ specialization.

F.7 Importer finance and take-or-pay exposure

Extend the model to include an importer-side additive payoff wedge $K_I \geq 0$ and a take-or-pay exposure term. Let $T(\phi, e)$ denote the importer's take-or-pay liability under green commitment, where $T(0, 0) = 0$. Alignment-contingent relief corresponds to $T_e(\phi, e) \leq 0$ locally.⁵⁶

In the linear–quadratic specialization, the commitment margin becomes

$$-De + N(\phi) + K_I + T(\phi, e) = 0,$$

so the threshold is characterized implicitly by this condition. Under a local affine approximation in effort,

$$T(\phi, e) \approx T(\phi, 0) + T_e(\phi, 0)e,$$

the modified threshold can be written in the same gap-over-slope form as

$$e_g^*(\phi) = \frac{N(\phi) + K_I + T(\phi, 0)}{D - T_e(\phi, 0)}, \quad (51)$$

provided $D - T_e(\phi, 0) > 0$.

Baseline take-or-pay exposure and importer finance enter the numerator and raise the threshold. Alignment-contingent relief corresponds to $T_e(\phi, 0) < 0$, which increases

⁵⁶If the intended interpretation is that greater incumbent flexibility mitigates take-or-pay exposure, then T is weakly decreasing in ϕ as well. If instead flexibility is purchased via tighter take-or-pay provisions, then T can be increasing in ϕ . The threshold results below do not require a sign on T_ϕ .

the effective slope $D - T_e(\phi, 0)$ and lowers the threshold by making exporter effort more effective at relaxing the commitment margin.

F.8 Interaction effects: technical-information and certification

Proposition 7 (Instrument interaction). *The cross-partial of the threshold with respect to baseline eligibility variance u_0 and baseline technical variance σ_0^2 satisfies*

$$\frac{\partial^2 e_g^*}{\partial \sigma_0^2 \partial u_0} = -\frac{\Pi_{u_0 \sigma_0^2}}{\Pi_e} + R, \quad (52)$$

where the leading term is $-\Pi_{u_0 \sigma_0^2}/\Pi_e = \Omega_{u \sigma^2}/\Pi_e$ when A does not depend directly on σ_0^2 , and

$$R = -\frac{\Pi_{u_0 e}}{\Pi_e} \frac{\partial e_g^*}{\partial \sigma_0^2} - \frac{\partial e_g^*}{\partial u_0} \cdot \frac{1}{\Pi_e} \left[\Pi_{e \sigma_0^2} + \Pi_{ee} \frac{\partial e_g^*}{\partial \sigma_0^2} \right] \quad (53)$$

collects terms involving curvature of Π in e and cross-derivatives with respect to e and the variance parameters.

Under the linear-quadratic specification in Appendix F, the commitment condition is linear in e and σ_0^2 and u_0 enter Ω additively. This implies $\Pi_{ee} = 0$, $\Pi_{u_0 e} = 0$, $\Pi_{e \sigma_0^2} = 0$, and $\Omega_{u \sigma^2} = 0$, so $R = 0$ and the cross-partial vanishes exactly:

$$\frac{\partial^2 e_g^*}{\partial \sigma_0^2 \partial u_0} = 0 \quad (LQ, \text{ exact}). \quad (54)$$

The effects of reducing eligibility variance and reducing technical variance on the commitment threshold are therefore exactly additive under the baseline LQ specification.

Under an extended specification in which the option value exhibits complementarity between technical and eligibility variance ($\Omega_{u \sigma^2} > 0$), the leading term $\Omega_{u \sigma^2}/\Pi_e$ is positive. The remainder R depends on curvature terms from two sources: the recognition-probability block (through $A_{u_0 e}$, which involves p_{ue}) and the option-value block (through Ω -derived contributions to $\Pi_{u_0 e}$, $\Pi_{e \sigma_0^2}$, and Π_{ee}). In the region where $\Omega_u > A_u$:

- (i) If $\Omega_{u \sigma^2} = 0$ and $R = 0$ (as in LQ), a reduction in σ_0^2 has no effect on the marginal threshold effect of reducing u_0 .
- (ii) If $\Omega_{u \sigma^2} > 0$ and $|R|$ is small relative to the leading term, a reduction in σ_0^2 lowers the marginal threshold reduction from lowering u_0 . The effects are sub-additive.

Proof. The threshold satisfies $\Pi(e_g^*; u_0, \sigma_0^2) = 0$. By the implicit function theorem, $\partial e_g^*/\partial u_0 = -\Pi_{u_0}/\Pi_e$, where $\Pi_e > 0$ by Assumption 7. Differentiating with respect to σ_0^2 via the quotient rule, and noting that both Π_{u_0} and Π_e are evaluated at $e = e_g^*(\sigma_0^2, u_0)$:

$$\frac{\partial^2 e_g^*}{\partial \sigma_0^2 \partial u_0} = -\frac{1}{\Pi_e} \left[\Pi_{u_0 \sigma_0^2} + \Pi_{u_0 e} \frac{\partial e_g^*}{\partial \sigma_0^2} \right] - \frac{\partial e_g^*}{\partial u_0} \cdot \frac{1}{\Pi_e} \left[\Pi_{e \sigma_0^2} + \Pi_{ee} \frac{\partial e_g^*}{\partial \sigma_0^2} \right]. \quad (55)$$

The first bracket collects the total derivative of Π_{u_0} with respect to σ_0^2 (direct cross-partial plus the indirect effect through the moving evaluation point e_g^*). The second bracket collects the analogous total derivative of Π_e , weighted by $-\partial e_g^*/\partial u_0 = \Pi_{u_0}/\Pi_e$.⁵⁷

⁵⁷The sign of the second bracket follows from the quotient rule applied to $-\Pi_{u_0}/\Pi_e$: the numerator derivative contributes the first bracket, while the denominator derivative contributes $+(\Pi_{u_0}/\Pi_e^2) \cdot d\Pi_e/d\sigma_0^2$. Since $\Pi_{u_0}/\Pi_e = -\partial e_g^*/\partial u_0$, this becomes $-(\partial e_g^*/\partial u_0)(1/\Pi_e) \cdot d\Pi_e/d\sigma_0^2$.

Under the maintained assumption that A does not depend directly on σ_0^2 , the direct cross-partial is $\Pi_{u_0\sigma_0^2} = -\Omega_{u\sigma^2}$. Collecting all terms beyond the leading $-\Pi_{u_0\sigma_0^2}/\Pi_e$ into R yields (52)–(53).

Under the LQ specification: the commitment condition $\Pi = -e D(\phi) + N(\phi, u_0, \sigma_0^2)$ is linear in e , so $\Pi_{ee} = 0$ exactly; the numerator N is additively separable in u_0 and σ_0^2 , so $\Pi_{u_0e} = 0$ and $\Pi_{e\sigma_0^2} = 0$; and σ_0^2 and u_0 enter Ω additively, so $\Omega_{u\sigma^2} = 0$. Therefore every term in (55) vanishes, establishing (54).

Under an extended specification with $\Omega_{u\sigma^2} > 0$ (for example, an interaction term $\beta_{12}\sigma^2 u$ in Ω), the leading term is $\Omega_{u\sigma^2}/\Pi_e > 0$. The remainder R involves curvature from two sources. On the recognition-probability side, Π_{u_0e} includes $A_{u_0e} = \bar{\tau}(\iota^{\text{def}} - \iota^g)p_{ue}$, which is small in the tail region where $|d|$ is large and p_{ue} is bounded. On the option-value side, the interaction term generates contributions of order $\beta_{12}\gamma$ to Π_{u_0e} , $\beta_{12}\eta$ to $\Pi_{e\sigma_0^2}$, and $\beta_{12}\gamma\eta$ to Π_{ee} . When the interaction coefficient β_{12} is modest relative to $\Omega_{u\sigma^2}$ —which holds when the complementarity in Ω is generated by a small perturbation away from the additive baseline—these option-value-side curvature terms are also small relative to the leading term, and the cross-partial has the sign of $\Omega_{u\sigma^2}$. $\square$

Consider two types of uncertainty-reducing interventions: a technical-information intervention that reduces both technical and eligibility variance by $(\Delta_\sigma, \Delta_u)$, and a certification intervention that further reduces eligibility variance by Δ_c . The total reduction in the threshold depends on the curvature of Ω .

Under the simple LQ specification with $\Omega = \beta_1\sigma^2 + \beta_2u$, the cross-partial $\Omega_{u\sigma^2} = 0$, and the effects are exactly additive:

$$\Delta e_g^*(\text{both}) = \Delta e_g^*(\text{technical-information}) + \Delta e_g^*(\text{certification}). \quad (56)$$

If Ω exhibits strict supermodularity in (σ^2, u) , so that $\Omega_{u\sigma^2} > 0$, then bundling the interventions yields diminishing returns relative to the sum of stand-alone moves. Reducing one source of uncertainty lowers the marginal value of reducing the other.

When reform capacity is limited, this logic suggests greater leverage from addressing the dominant source of uncertainty first. If technical uncertainty is the larger contributor to delay, front-end engineering and data transparency yield the larger initial threshold reduction before certification reform is pursued. If eligibility uncertainty is the larger contributor, certification clarity should come first.

G Exporter outside option: static analysis

This appendix analyzes how exporter-side alternative-importer depth affects the timing equilibrium. The analysis is static and does not address dynamic reputational or geopolitical effects.

G.1 Setup

Let $\psi \geq 0$ index the depth of the exporter’s alternative-importer option, interpreted as the credibility and value of redirecting capacity to other importers. The parameter ψ may shift both importer-side and exporter-side primitives.

Fix the incumbent flexibility level ϕ . Let $\Pi(e; \phi, \psi)$ denote the importer’s net gain from committing at $t = 0$ given exporter effort e . Assumption 7 implies $\Pi_e(e; \phi, \psi) > 0$ at the threshold.

On the exporter side, let $F_E(e; \psi)$ denote the exporter's gross benefit from inducing commitment at effort e , let $\Lambda_E(\psi)$ denote any additive exporter payoff shifter associated with the alternative-importer option, and let $C_E(e; \psi)$ denote the exporter's effort cost. Assume $C_E(0; \psi) = 0$, $C_{E,e} > 0$, and $C_{E,ee} > 0$.

The directions of these effects depend on how the alternative-importer move is implemented. A queue-based scarcity signal that raises the importer's cost of waiting increases the net gain from committing and tends to make Π_ψ positive. A move that raises perceived policy or eligibility risk reduces the net gain from committing and tends to make Π_ψ negative. A move that undermines cost-sharing arrangements raises the exporter's marginal cost of effort and tends to increase $C_{E,\psi}$.

G.2 Effect on the threshold

Define the threshold $e_g^*(\psi)$ implicitly by the commitment condition

$$\Pi(e_g^*(\psi); \phi, \psi) = 0. \quad (57)$$

Implicit differentiation gives

$$\frac{\partial e_g^*}{\partial \psi} = -\frac{\Pi_\psi(e_g^*(\psi); \phi, \psi)}{\Pi_e(e_g^*(\psi); \phi, \psi)}. \quad (58)$$

Since $\Pi_e > 0$, the move is commitment-friendly (lowers the threshold) if and only if

$$\Pi_\psi(e_g^*(\psi); \phi, \psi) > 0, \quad (59)$$

and it is commitment-hostile if $\Pi_\psi(e_g^*(\psi); \phi, \psi) < 0$. Queue and priority effects that raise the cost of waiting increase Π and tend to lower the threshold. Risk and eligibility effects that reduce the net gain from commitment decrease Π and tend to raise the threshold.

G.3 Effect on exporter payoff

Let the exporter's net benefit from inducing commitment be

$$\Delta_E(\psi) = F_E(e_g^*(\psi); \psi) + \Lambda_E(\psi) - C_E(e_g^*(\psi); \psi). \quad (60)$$

The total derivative is

$$\frac{d\Delta_E}{d\psi} = F_{E,\psi} + F_{E,e} \frac{\partial e_g^*}{\partial \psi} + \Lambda_{E,\psi} - C_{E,\psi} - C_{E,e} \frac{\partial e_g^*}{\partial \psi}, \quad (61)$$

with all terms evaluated at $(e_g^*(\psi), \psi)$. The move is exporter-rational if $d\Delta_E/d\psi \geq 0$.

G.4 Feasibility

Even if the move passes the commitment and return-on-investment tests, it must not push the threshold above the feasible effort bound: $e_g^*(\psi) \leq 1$. If the threshold exceeds unity, commitment becomes infeasible regardless of its effects on Δ_E .

G.5 Static threat screen

Combining these conditions yields a static threat screen. A marginal increase in alternative-importer depth is value-creating only if three conditions hold simultaneously: the commitment test is satisfied ($\partial e_g^*/\partial \psi \leq 0$), the return-on-investment test is satisfied ($d\Delta_E/d\psi \geq 0$), and feasibility is maintained ($e_g^*(\psi) \leq 1$).

Moves that primarily operate through the queue/priority channel, raising foregone rents without increasing perceived risk, are most likely to pass all three tests. Moves that primarily raise eligibility variance or undermine cost sharing are likely to fail.

H Carbon price and foregone rents

The main text (Proposition 2) decomposes the carbon-price effect on the threshold into an expected-assigned-intensity term and a foregone-rent channel. Appendix F reports the corresponding LQ partial derivative, $\partial e_g^*/\partial \bar{\tau} = \iota^g/D$, which holds foregone rents fixed. This appendix derives the total effect when foregone rents vary with the carbon price level and clarifies its relationship to the standard Pigouvian substitution effect.

H.1 The foregone-rent channel

In the main text, the foregone first-mover rent $\delta_{I,0}^g$ is treated as independent of the carbon price level $\bar{\tau}$. A fuller specification allows $\delta_{I,0}^g$ to depend on $\bar{\tau}$ because higher carbon prices raise the cost of continuing on fossil during the waiting period. The maintained sign restriction is

$$\delta_{I,0,\bar{\tau}}^g(\bar{\tau}) > 0.$$

H.2 General sign-based characterization

Allow $\delta_{I,0}^g$ to depend on $\bar{\tau}$. The commitment condition is $\Pi(e; \phi, \bar{\tau}) = 0$. Applying the implicit function theorem gives

$$\frac{de_g^*}{d\bar{\tau}} = -\frac{\Pi_{\bar{\tau}}}{\Pi_e}. \quad (62)$$

Under Assumption 7, $\Pi_e > 0$ at the threshold. The carbon price enters the commitment margin directly through the assigned intensity term and indirectly through foregone rents, so

$$\Pi_{\bar{\tau}} = \iota^g - \delta_{I,0,\bar{\tau}}^g(\bar{\tau}). \quad (63)$$

If the foregone-rent channel dominates the direct green-charge term, $\delta_{I,0,\bar{\tau}}^g(\bar{\tau}) > \iota^g$, then $\Pi_{\bar{\tau}} < 0$ and $de_g^*/d\bar{\tau} < 0$: higher carbon prices lower the threshold and accelerate commitment.

H.3 Local-slope comparison

Define the local foregone-rent sensitivity $\theta_{\tau}(\bar{\tau}) \equiv \delta_{I,0,\bar{\tau}}^g(\bar{\tau})$. Using (63) and $\Pi_e = D > 0$ at the threshold yields

$$\frac{de_g^*}{d\bar{\tau}} = \frac{\iota^g - \theta_{\tau}(\bar{\tau})}{D}. \quad (64)$$

For certified goods with $\iota^g \approx 0$ and $\theta_{\tau}(\bar{\tau}) > 0$, the total effect is negative: higher carbon prices lower the threshold through the fossil-side foregone-rent channel. The partial

derivative reported in the main text, $\partial e_g^*/\partial \bar{\tau} = \iota^g/D$, isolates the direct green-charge term by holding $\delta_{I,0}^g$ fixed.

A magnitude comparison on the timing margin can be written using (64). Under the interaction form, the eligibility-variance sensitivity is $\partial e_g^*/\partial u_0 = (\beta_2 + \xi_{\Omega,1}\phi)/D(\phi)$ and the total carbon-price sensitivity is $|de_g^*/d\bar{\tau}| = |\iota^g - \theta_\tau(\bar{\tau})|/D(\phi)$. The common denominator cancels in the ratio of local slopes, but the ratio depends on the normalization of u because rescaling u changes the variance-sensitivity coefficient. The calibration in Section 6 therefore reports an equivalent carbon-price change $\Delta \bar{\tau}^{\text{eq}}$ that is invariant to rescaling of u : the product of the variance-sensitivity coefficient and the MRA-induced variance shift does not depend on the units in which eligibility variance is measured.

H.4 Interpretation

The analysis in this appendix clarifies three points.

First, the partial derivative $\partial e_g^*/\partial \bar{\tau} = \iota^g/D$ in the main text isolates the certification margin by holding foregone rents fixed. That object captures how a higher carbon price affects timing through the assigned-intensity term in the importer's commitment condition. It does not claim that carbon pricing is irrelevant to investment timing in general.

Second, once foregone rents are allowed to depend on the carbon price, the total effect of carbon pricing on timing can have the standard Pigouvian sign. In particular, if the foregone-rent sensitivity $\theta_\tau(\bar{\tau}) \equiv \delta_{I,0,\bar{\tau}}^g(\bar{\tau})$ exceeds the assigned intensity term ι^g , then $de_g^*/d\bar{\tau} < 0$: higher carbon prices lower the threshold and accelerate commitment by increasing the cost of continued fossil use during the waiting period. The main text emphasizes the certification margin because that margin is where the model adds structure relative to standard fossil-side substitution logic.

Third, the eligibility-variance-dominance result is best interpreted as a comparison of magnitudes on the timing margin. For certified goods with near-zero assigned intensity, the direct carbon-charge channel is weak ($\iota^g \approx 0$). Carbon pricing can still matter through the foregone-rent channel, summarized by $\theta_\tau(\bar{\tau})$. The option-value channel driven by eligibility variance can nonetheless be quantitatively larger. In the local-slope representation, eligibility variance dominates carbon pricing on the timing margin when $\beta_2 > \theta_\tau(\bar{\tau})$. The policy implication is that certification clarity is a lever alongside carbon price support, not a substitute for it.

I Empirical detail

This appendix provides additional detail on the data, measurement, and regression specifications underlying Section 7.

I.1 Data and sample

The IEA Hydrogen Production Projects Database (September 2025) records 2,625 hydrogen production projects worldwide.⁵⁸ The analysis excludes 153 Chinese projects, leaving 2,463 projects. The pipeline sample used in descriptive tables restricts to projects at Concept, Feasibility Study, or FID/Construction stages ($N = 1,858$), of which 186 have

⁵⁸After dropping eight aggregate rows without project-level identifiers and one duplicate reference number.

Table 8: Concordance: Hand-Coded Export Flag and Ammonia Proxy

	Ammonia = 0	Ammonia = 1	Total
Hand-coded domestic	165 (31.3%)	64 (12.1%)	229
Hand-coded export	90 (17.0%)	209 (39.6%)	299
Total	255	273	528

Notes: Cell percentages in parentheses. Agreement rate: 70.8%.

Table 9: FID Rates Under Hand-Coded Export Classification (Pipeline Projects)

	Domestic	Export	Difference	<i>p</i> -value
Electrolysis	4.7% (<i>N</i> = 106)	1.6% (<i>N</i> = 258)	+3.2 pp	0.077
CCUS	10.1% (<i>N</i> = 99)	9.8% (<i>N</i> = 41)	+0.3 pp	1.000

Notes: Pipeline projects with hand-coded export classification. *p*-values from Pearson χ^2 test (electrolysis) and Fisher’s exact test (CCUS).

reached FID by September 2025. The logit sample further restricts to electrolysis and CCUS technologies with non-missing capacity (*N* = 1,608, 170 FID events).

I.2 Export classification

The ammonia product listing serves as a proxy for export orientation in the full IEA sample. Among 528 projects where both the ammonia proxy and hand-coded classification are available, agreement is 70.8 percent. Table 8 reports the concordance. Among hand-coded export projects, 30 percent lack an ammonia listing, indicating that some export projects use carriers other than ammonia or do not list their product. Among hand-coded domestic projects, 28 percent list ammonia, reflecting domestic fertilizer production.

For CCUS projects, the hand-coded classification covers all 140 pipeline projects (99 domestic, 41 export). None of the 79 CCUS projects added in the hand-coding expansion are export-oriented; they consist of refinery carbon capture, fertilizer-plant carbon capture, industrial decarbonization, and waste-to-hydrogen projects located primarily in the United Kingdom, the Netherlands, the United States, Canada, and South Korea. All CCUS export projects identified in the hand-coding had already been classified in the earlier coding round. The hand-coding expansion did not identify additional export-oriented CCUS projects, which reduces concern that the near-zero clean-coded CCUS gap is driven by omitted export cases.

I.3 Hand-coded FID rates

Table 9 reports FID rates under the hand-coded export classification for the pipeline subsample where both measures are available.

The electrolysis export penalty survives cleaner classification: export projects reach FID at roughly one-third the domestic rate under both measures. The CCUS contrast reverses: under the ammonia proxy, export CCUS appears advantaged (15.6 versus 7.4 percent); under clean coding, the gap is essentially zero (9.8 versus 10.1 percent). The

Table 10: Logit: FID Status, Ammonia Proxy (Comparison)

	(1)
Export (ammonia)	1.20* (0.69)
Electrolysis	−0.64 (0.48)
Export × Electrolysis	−1.80** (0.76)
Log capacity	−0.49*** (0.04)
<i>N</i>	1,608
FID events	170

Notes: Same specification and sample as Table 7, but using ammonia product listing to identify export orientation for all technologies. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

CCUS reversal in Table 6 is therefore best interpreted as a proxy artifact arising from the misclassification of domestic ammonia-for-fertilizer projects.

I.4 Ammonia proxy logit

Table 10 reports the same logit specification as Table 7 but using the ammonia proxy for all technologies rather than the corrected classification. The electrolysis-specific export effect is virtually unchanged (−0.60 versus −0.61 in the corrected specification). The CCUS export coefficient is larger (+1.20, $p = 0.081$) and the interaction term is more sharply negative (−1.80, $p = 0.017$), both consistent with the ammonia proxy overstating the CCUS export FID rate.

I.5 Hand-coded subsample logit

Table 11 reports the logit using the hand-coded export flag for both technologies, restricted to the 459 pipeline projects with non-missing capacity and hand-coded classification. Only 20 FID events are available in this sample. The electrolysis export effect is −1.00 ($p = 0.147$): the sign is preserved but precision is limited by the small number of events.

I.6 Capacity missingness

Of the 140 hand-coded CCUS pipeline projects, 45 (32 percent) have missing capacity data and are excluded from the logit sample. The missing observations are disproportionately domestic: 36 of 45 missing-capacity CCUS projects (80 percent) are domestic, compared to 63 of 95 non-missing projects (66 percent). This compositional difference means that the logit sample overrepresents export CCUS relative to the rate-table sample, which may contribute to the positive CCUS export point estimate in the logit despite the near-zero raw-rate gap under clean coding.

Table 11: Logit: FID Status, Hand-Coded Subsample

	(1)
Export (hand-coded)	0.31 (0.70)
Electrolysis	−0.89 (0.61)
Export × Electrolysis	−1.32 (0.96)
Log capacity	−0.16 (0.15)
N	459
FID events	20

Notes: Pipeline projects with hand-coded export classification and non-missing capacity. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

I.7 Alternative outcome

Table 12 reports the logit using an alternative outcome that counts both FID/Construction and Operational projects as having advanced. The sample adds Operational projects to the pipeline ($N = 1,921$, 483 events). The electrolysis export effect is -0.25 ($p = 0.379$) and the interaction is -1.43 ($p = 0.033$). The weaker electrolysis export effect is expected: Operational projects are disproportionately older vintage, and the certification uncertainty relevant to the theory is a more recent phenomenon.

I.8 Alternative explanations

Several alternative explanations for the technology-specific export penalty deserve consideration.

Export electrolysis projects may differ systematically from domestic projects in scale ambitions or developer characteristics. Projects targeting export markets are often larger and backed by consortia with longer decision horizons, which could independently delay FID. The logit controls for log capacity but not for unobserved heterogeneity in developer type, financing structure, or host-country regulatory conditions.

CCUS-based export projects may benefit from geographic concentration in jurisdictions with established regulatory frameworks, notably the United States where the 45Q tax credit provides a recognized certification pathway, or from importer-country support programs that effectively pre-certify output. If so, the near-zero CCUS export penalty reflects favorable policy environments rather than inherently lower eligibility variance.

These alternatives cannot be ruled out with the available data. The evidence presented in Section 7 is a descriptive plausibility check, not a causal test. The technology-specific pattern is nevertheless difficult to reconcile with explanations that predict a uniform export penalty across technologies, because such explanations cannot account for the absence of a comparable CCUS gap.

Table 12: Logit: FID or Operational Status (Alternative Outcome)

	(1)
Export (ammonia)	1.19* (0.62)
Electrolysis	-1.37*** (0.40)
Export $\times$ Electrolysis	-1.43** (0.67)
Log capacity	-0.77*** (0.04)
N	1,921
FID + Operational events	483

Notes: Dependent variable equals one if project status is FID/Construction or Operational. Sample includes pipeline projects and Operational projects, excluding China, with non-missing capacity. Export identified by ammonia product listing. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.