Voting Rules and the Price of Peace[*]

Kenneth Khoo[†]
kenneth.khoo@nus.edu.sg

Abstract

I integrate ex ante shareholder voting with ex post litigation/ratification in a unified model of controlling shareholder transactions. A controller who exaggerates deal quality to win minority approval must later buy peace if shareholders sue, through a ratification vote whose cost increases in the size of the exaggeration, the toughness of the ratification threshold, and the strictness of judicial scrutiny. When the same threshold governs both approval and ratification, tougher votes simultaneously raise the bar for persuasion and the price of peace, reversing the standard prediction that stricter voting rules may increase misrepresentation. The reversal arises because the approval threshold does two jobs at once in the coupled regime: it sets the standard for persuasion and the cost of peace, and the enforcement channel can dominate the mechanical expansion of the pooling region. The welfare analysis decomposes any governance regime's shortfall into false positives, false negatives, and enforcement costs, and identifies four institutionally interpretable parameters that jointly determine whether coupling is beneficial. Mapping these parameters onto ten jurisdictions shows that coupling improves welfare in tunneling-prone markets but harms it where over-deterrence is the binding constraint. The model further proves that independent committees and tough ratification standards are procedural complements, not substitutes: the sequential structure of governance generates this complementarity endogenously, because the cost of committee independence is invariant to enforcement toughness while the benefit scales with it. Either/or safe harbors like Delaware's 2025 Senate Bill 21 are therefore generically welfare-inferior to joint requirements.

[*]I am grateful to Deborah Beim, Lucian Bebchuk, Albert Choi, Xinyu Hua, Lewis Kornhauser, Kathryn Spier, Alan Miller, and other participants at the Fourteenth Annual Law and Economic Theory Conference (LET XIV) (Harvard Law School) for helpful comments, and to Roberto Tallarita for invaluable discussions.

[†]National University of Singapore, Faculty of Law.

1 Introduction

Controlling shareholders are the dominant ownership form in global capital markets. La Porta et al. (1999) documented that in two-thirds of large firms across 27 wealthy economies, a single shareholder controls more than 20% of votes. Claessens et al. (2000) found that in East Asia, a single family controls more than 40% of listed firms. Even in the United States, where dispersed ownership is the textbook norm, dual-class structures, family control, and private-equity-backed public companies make controlling shareholders increasingly prevalent (Bebchuk and Kastiel, 2017). The transactions these controllers initiate—freeze-outs, related-party deals, going-private mergers—collectively transfer enormous value across global capital markets, making the governance of controller conflicts one of the central problems in corporate law.

The central governance challenge posed by these transactions is that the controller sits on both sides of the deal. The controller has private information about deal quality θ and stands to extract private benefits $B(\theta)$ from implementation. This creates an adverse-selection problem: the controller knows whether a proposed transaction creates or destroys value for minority shareholders, and has incentives to push through value-destroying deals that generate private benefits. The fundamental design question is how to construct governance mechanisms that screen out bad deals while allowing good deals to proceed.

The conventional response relies on shareholder voting, board independence, and judicial review. A natural conjecture is that stricter voting rules should reduce misrepresentation: raise the bar, and some controllers who would have exaggerated will find it too costly and drop out. That conjecture is incomplete. When the approval threshold rises, controllers whose deals are genuinely bad do drop out, because exaggeration across a wider gap has become prohibitively expensive. But the same threshold increase sweeps in controllers whose deals are almost good enough but now fall just short of the new, higher bar. These newly displaced controllers need only a small exaggeration to clear it, and because their deals are nearly value-creating, the benefit of approval easily justifies that modest cost. More light exaggerators enter at the top than heavy exaggerators exit at the bottom, so total misrepresentation rises. This is not a pathological corner case; it is the generic one-stage prediction under standard distributional assumptions.

Empirical evidence is consistent with this threshold-bunching logic. Bach and Metzger (2019) document a density discontinuity at the majority threshold for U.S. governance proposals; Burgstahler and Dichev (1997) show that reported earnings pile up just above zero.[1] Restrepo (2021) finds that after *MFW* incentivized dual protections, MoM conditions rose sharply yet deal premiums showed no significant change—consistent with reshuffling rather than reducing pooling. The question is what breaks this logic; the answer is the enforcement channel that operates after the vote.

This paper studies conflicted-controller transactions as a two-stage enforcement problem. In

[1]The bunching methodology is developed in Saez (2010) and surveyed by Kleven (2016); see also Ewens et al. (2024) for bunching estimation at governance-relevant regulatory thresholds. Cross-jurisdictional evidence: Fried et al. (2020) (Israel); Li (2021) (India); Becht et al. (2016) (UK); Holderness (2018) (cross-country); Boone et al. (2018) (Delaware).

Stage 1, the controller seeks minority approval under a voting threshold π. In Stage 2, if the deal is challenged, the controller tries to extinguish claims through settlement and ratification, subject to threshold r and judicial willingness to credit the cleansing vote. The key state variable is the disclosure gap $g := m - \theta$: larger gaps make litigation more likely and make court-creditable peace more expensive. That Stage-2 technology feeds back into Stage-1 disclosure incentives. The central mechanism is procedural complementarity: a stricter approval rule raises the ex ante bar to pass, and if ratification thresholds are coupled to approval thresholds ($r(\pi) = \pi$), the same reform also hardens the ex post buy-peace technology. Higher π raises the voting bar $u(\pi)$; if $r'(\pi) > 0$, it also raises expected Stage-2 enforcement costs for any gap g; and anticipating this, controllers reduce near-bar exaggeration, so pooling can fall as π rises. The contribution is therefore not "which threshold is best" in isolation, but how approval and cleansing rules interact through enforcement to shape ex ante disclosure.

The paper develops this mechanism in four steps. First, it formalizes a reduced-form enforcement technology $L(g, r, \sigma)$ grounded in litigation and ratification, derives the equilibrium at a fixed voting bar, and shows that coupling $r(\pi)$ can reverse the benchmark prediction that tighter votes increase pooling—a reversal formalized by the litigation-adjusted density-weighted boundary responsiveness condition (LA-DWBR). Second, it develops a welfare decomposition identifying four institutionally interpretable parameters—the deadweight share of enforcement costs (ω), the convexity of those costs in the misrepresentation gap (κ), the sensitivity of enforcement to the ratification threshold (L_r), and the quality of the marginal deterred deal (θ^*)—and maps them onto ten jurisdictions: coupling is welfare-improving where marginal deals are value-destroying (India, Hong Kong, China) and welfare-harmful where they are value-creating (Japan, the UK), with Delaware as a contested case where three of four parameters favor coupling despite S.B. 21's move toward decoupling. Third, it establishes that independent committees and tough ratification standards are procedural complements—a formal supermodularity result implying that either/or safe harbors of the kind S.B. 21 introduced are generically welfare-inferior to joint requirements. Fourth, it shows that private ordering fails where coupling is most needed: controllers choosing their own governance packages internalize minority-value gains only in proportion to their ownership stake, producing a governance externality that is most severe in high-tunneling, concentrated-ownership markets.

The remainder of the paper proceeds as follows. Section 2 maps the model to doctrine and related literature. Section 3 presents the core two-stage model. Section 4 develops the Stage-2 enforcement and ratification technology. Section 5 characterizes Stage-1 equilibrium and the litigation-adjusted comparative statics under coupled versus decoupled rules. Section 6 develops the welfare analysis, comparative institutions, and extensions on private ordering and committee design. Section 7 concludes. The Appendix collects deferred proofs and microfoundation details; the Online Appendix contains robustness extensions.

2 Institutional Mapping and Related Literature

This paper maps directly to the legal sequence in controller transactions. Stage 1 is minority approval. Stage 2 is litigation, settlement, and potential ratification/cleansing, with judicial screening of whether votes were fully informed, uncoerced, and independent. The parameter σ captures the risk that courts refuse to credit purchased or tainted votes, and r captures how demanding ratification is. Together they determine how expensive it is to buy peace after approval.

Under the *MFW* framework (*Kahn v. M&F Worldwide Corp.*, 2014), a controlling shareholder seeking business judgment review of a freeze-out merger must satisfy two conditions: negotiation by an empowered, independent special committee, and approval by a fully informed, uncoerced majority of the minority shareholders. Satisfying both shifts the standard of review from entire fairness to business judgment; failing either leaves the controller exposed to entire-fairness scrutiny, with the attendant litigation risk and potential rescissory damages (Fiegenbaum, 2016, 2019; Subramanian, 2007; Pritchard, 2004; Lazarus and McCartney, 2011). In the model, MFW implements the coupled regime: $r = \pi$, with high committee independence ι and high ratification toughness r required jointly.

Corwin v. KKR Financial Holdings (2015) extended a version of this logic to non-controller transactions: a fully informed, uncoerced stockholder vote can "cleanse" a board-level conflict and restore business judgment review (Hill and McDonnell, 2011). The Corwin doctrine formalizes the Stage 2 mechanism—the possibility that a post-approval vote extinguishes claims—while simultaneously illustrating its fragility: courts have declined to apply Corwin where disclosure was inadequate or where structural coercion was present (the σ parameter) (Restrepo and Subramanian, 2015; Licht, 2020; Cox et al., 2019; Gevurtz, 2019). Recent DGCL Section 144 amendments (SB 21, March 2025) expanded the statutory menu by codifying safe-harbor pathways that permit either/or procedural protections for non-freeze-out interested transactions—the decoupled corner solution that the model's supermodularity result predicts is welfare-inferior.

The coupling extends beyond Delaware. In the UK, special resolutions require 75% of eligible votes for both approval and ratification.[2] This is formally coupled ($r(\pi) = \pi$), but effective stringency depends on whether the controller's shares count toward the denominator. A 60%-controller needs only 15 percentage points of independent support—unlike MoM, where the controller is excluded. The model predicts coupling disciplines misrepresentation only when both thresholds are genuinely demanding from the minority's perspective. At the extreme, the 90% squeeze-out threshold in the UK and Hong Kong requires 30 percentage points of independent support, where the ratification channel operates with maximal force.[3]

Related work falls into three groups. First, voting-threshold and signaling models provide the benchmark bar-boundary logic (Kartik, 2009; Bebchuk and Kahan, 2000; At et al., 2015; Levy, 2005; Dalkır et al., 2019); this paper endogenizes the signaling cost through Stage 2 enforcement, so the cost surface responds to the screening rule. Second, controller-transaction scholarship documents

[2] Section 283, UK Companies Act 2006. Australia similarly requires 75% supermajorities (Section 9 with Section 250MA, Corporations Act 2001).

[3] Section 979, UK Companies Act 2006; Part 13, Hong Kong Companies Ordinance.

institutional variation (inter alia Gilson and Gordon, 2003; Gilson, 2006a,b; Djankov et al., 2008; Dyck and Zingales, 2004; Goshen, 2003; Goshen and Hamdani, 2019; Gilson and Schwartz, 2013; Gutiérrez Urtiaga and Sáez Lacave, 2017; Enriques and Tröger, 2019; Dammann, 2019; Pacces, 2019; Fiegenbaum, 2019; Lovo et al., 2021; Lauterbach et al., 2021; Bolton et al., 2003) but does not formalize the feedback from post-vote enforcement into ex ante disclosure. Third, the litigation and settlement literature (Bebchuk, 1984) explains why enforcement costs contain both transfer and deadweight components. The model integrates these strands: ex post cleansing technology changes ex ante disclosure, with the welfare cost governed by ω.

The one-stage threshold-bunching model is kept as a benchmark foil (see Appendix B). The paper's identity is the two-stage enforcement feedback under coupling $r(\pi)$.

3 Model

This section presents the core model. Subsequent results cite only the numbered statements below.

3.1 Timeline (Stages 0–2)

The game proceeds in three stages:

Stage 0. Nature. Project quality θ is drawn from F on $[\underline{\theta}_0, \overline{\theta}_0]$ and observed privately by the controlling shareholder (CS).

Stage 1. Disclosure & Vote. The CS sends a public message $m \in \mathbb{R}$. Minority shareholders update beliefs and vote under the prevailing approval rule (MoM or supermajority with quota π). If the proposal fails, the game ends with payoff zero.

Stage 2. Extraction, Litigation & Ratification. Conditional on approval:

(a) The CS extracts private benefits $B(\theta)$ (or, in the continuous-extraction extension, chooses $x \in [0, B(\theta)]$).

(b) Minority shareholders may file suit. The probability and expected cost of litigation are increasing in the *misrepresentation gap* $g := m - \theta$ (or $g := m^{\text{req}}(u) - \theta$).

(c) The CS may attempt to extinguish claims by buying support for a *ratification/settlement vote* requiring threshold r. Courts may decline to treat a stockholder vote as cleansing if it is not fully informed and uncoerced; in the model, this is captured by probability $\sigma \in [0, 1)$ that the court refuses to credit the ratification vote.

(d) The litigation/ratification subgame resolves, producing a reduced-form expected enforcement cost $L(g, r, \sigma)$.

Solving backward from Stage 2, the enforcement cost $L(g, r, \sigma)$ enters the CS's Stage-1 problem exactly as a signaling cost would, but with the additional feature that L depends on the ratification threshold r and the court-independence parameter σ—which, in the carry-forward baseline, includes $r = \pi$ (the same quota as the approval rule).

The timeline is solved in reverse: the controller, contemplating Stage 1 disclosure, anticipates Stage 2 enforcement exposure. This forward-looking calculation transforms institutional parameters—r, σ, shareholder resistance—into a disciplinary force on Stage 1 behavior.

The carry-forward baseline $r(\pi) = \pi$ is grounded in practice. In Delaware post-MFW, the same conditions define cleansing at both stages; in the UK, the same 75% supermajority governs both the original resolution and any ratification (Companies Act 2006, s.283). When r moves with π, this feedback creates a disciplinary channel absent from one-stage models.

Unless otherwise stated, the model uses

$$\alpha \in (0,1), \quad \pi \in [1/2, 1), \quad r \in (0,1], \quad \sigma \in [0,1), \quad \lambda = 1/N \in (0,1),$$

with $\theta \in [\underline{\theta}_0, \overline{\theta}_0]$ and message $m \in \mathbb{R}$. The misrepresentation gap is

$$g := m - \theta,$$

and at the boundary we write $g^* := m^{\text{req}}(u) - \theta^*$. In the ratification-vote microfoundation, recovery satisfies $R \in [0,1]$ and cleansing is feasible only if

$$0 < r \leq 1 - \sigma.$$

For $r > 1 - \sigma$, the microfounded cleansing cost is unbounded and those points are excluded from comparative-statics corridors. Accordingly, whenever the paper specializes to that microfoundation and differentiates with respect to r or σ, analysis is restricted to the feasible interior region.

3.2 Controlling Shareholders

3.2.1 Project Quality

There exists a single *controlling shareholder* (CS) who permanently owns a strictly positive but non-unanimous voting stake $\alpha \in (0,1)$, while a large but finite number N of symmetric minority shareholders–which will be discussed later–hold the residual $1 - \alpha$.

The quality of a project is denoted by the variable θ. It is realised once, observed privately by the CS, and then remains fixed throughout the game. Importantly, θ is distributed $\theta \sim F$, and F is continuous on the compact interval $\left[\underline{\theta}_0, \overline{\theta}_0\right] \subset \mathbb{R}$, where the support endpoints satisfy

$$[\,\underline{\theta}_0,\ \overline{\theta}_0], \qquad \underline{\theta}_0 < 0 < \overline{\theta}_0. \tag{P1}$$

We assume the following:

Assumption 1 (Log-concavity and monotone shoulders (P2))**.** *The type distribution F has a den-*

sity $f := F'$ that is strictly positive on $[\underline{\theta}_0, \overline{\theta}_0]$ and log-concave*:*

$$f(x) > 0 \quad \textit{and} \quad \log f(\cdot) \textit{ is concave on } [\underline{\theta}_0, \overline{\theta}_0].$$

Let $u_{\text{MoM}} \in (\underline{\theta}_0, \overline{\theta}_0)$ denote the (unique) mode of f. Then f has monotone shoulders*: it is* nondecreasing *on $[\underline{\theta}_0,\ u_{\text{MoM}}]$ (left shoulder) and* nonincreasing *on $[\ u_{\text{MoM}},\ \overline{\theta}_0]$ (right shoulder), where the belief-cutoff u_{MoM} is defined later in V3 (Section 5.2.3).*

3.2.2 Private Benefits (pure perk + diversion)

Implementing the project gives the controller private benefits $B(\theta)$ on top of the pro-rata payoff $\alpha\theta$. We explicitly decompose

$$B(\theta) \ = \ b_0 \ + \ D(\theta),$$

where $b_0 \in \mathbb{R}$ is a *pure perk* (does not reduce minorities' payoff) and $D(\theta)$ is the diversionary component that does. We assume

$$D \ : \ [\,\underline{\theta}_0, \overline{\theta}_0] \longrightarrow \mathbb{R}, \qquad D \in C^1, \qquad 0 < B'(\theta) < 1 - \alpha \quad \big(\forall\, \theta\big), \tag{P3}$$

and normalize $D(\underline{\theta}_0) = 0$. Because b_0 is constant, without loss of generality we have $B'(\theta) = D'(\theta)$ on $[\underline{\theta}_0, \overline{\theta}_0]$. The controller's prize uses the full $B(\theta) = b_0 + D(\theta)$, while minorities subtract only $D(\theta)$. Hence the net minority payoff

$$(1 - \alpha)\,\theta \ - \ D(\theta)$$

is strictly increasing on $[\underline{\theta}_0, \overline{\theta}_0]$ by (P3), a property used in the single-crossing arguments that follow. If $B'(\theta) > 1-\alpha$ (the diversion–dominant case, deferred from the main text), the acceptance criterion reverses.

3.2.3 Signalling technology and enforcement cost

After observing θ, the controller sends a publicly observable message $m \in \mathbb{R}$. Disclosure itself is *costless ex ante*; however, any misrepresentation exposes the controller to expected Stage-2 enforcement costs (litigation, settlement, or failed ratification). As shown in Section 4 below, solving the Stage-2 subgame yields a reduced-form *enforcement cost* $C(m, \theta) \equiv L(g, r, \sigma)$ that depends on the misrepresentation gap $g := m - \theta$, the ratification threshold r, and the court-independence parameter σ.

For notation, we occasionally use $\Delta := m - \theta$ inside wedge-by-wedge derivative formulas. Thus g and Δ denote the same object.

For continuity with the existing proofs, we adopt the same smooth, convex, gap-based specification as before (cf. the lying-cost formulation in Kartik, 2009), but now interpret the parameters $\eta > 0$ and $\kappa \geq 0$ as enforcement-cost primitives that are *increasing in the ratification thresh-*

old r and in the *court-independence parameter* σ (see Assumption 2 below). Specifically, for all $(m, \theta) \in \mathbb{R}^2$:

$$C(m,\theta) \;\equiv\; L(g,r,\sigma) \;=\; \eta(r,\sigma)\,|m-\theta| \;+\; \frac{\kappa(r,\sigma)}{2}\,(m-\theta)^2, \qquad \text{with } C(\theta,\theta)=0. \tag{S1}$$

Equivalently, writing $\Delta := m - \theta$,

$$C(m,\theta) = \begin{cases} \eta(r,\sigma)\,(\theta-m) \;+\; \dfrac{\kappa(r,\sigma)}{2}\,(m-\theta)^2, & m<\theta, \\ 0, & m=\theta, \\ \eta(r,\sigma)\,(m-\theta) \;+\; \dfrac{\kappa(r,\sigma)}{2}\,(m-\theta)^2, & m>\theta. \end{cases} \tag{S2}$$

When r and σ are held fixed, all subsequent results reduce to the original signaling-cost model. The novel channel arises when $r = r(\pi)$ (carry-forward baseline), so that tightening π simultaneously raises $\eta(r, \sigma)$ and $\kappa(r, \sigma)$.

S3: Convexity and symmetry: C is convex in the gap Δ (sum of the convex map $\Delta \mapsto |\Delta|$ and the quadratic $\Delta \mapsto \frac{\kappa}{2}\Delta^2$) and *symmetric* in the sense $C(\theta + \Delta, \theta) = C(\theta - \Delta, \theta)$.

S4: Differentiability and marginal costs: C is C^∞ on each wedge $\{\Delta \neq 0\}$ and has a (single) kink at $\Delta = 0$. The partial derivatives with respect to the message and the type exist for $\Delta \neq 0$ and satisfy

$$C_m(m,\theta) = \begin{cases} -\eta + \kappa\Delta, & \Delta<0, \\ \eta + \kappa\Delta, & \Delta>0, \end{cases} \qquad C_\theta(m,\theta) = -\,C_m(m,\theta) = \begin{cases} \eta - \kappa\Delta, & \Delta<0, \\ -\eta - \kappa\Delta, & \Delta>0, \end{cases} \tag{S4}$$

with subgradient $C_m(0) \in [-\eta, \eta]$ at $\Delta = 0$. In particular,

$$C_m(0^-) = -\eta, \qquad C_m(0^+) = +\eta,$$

so both the up- and down-wedge display a strictly positive *linear lip* of size η at the truth.

S5: Monotonicity on (convex) wedges: On the up-wedge ($\Delta > 0$), the marginal signaling cost is strictly positive and strictly increasing in the size of the lie:

$$C_m(\Delta) = \eta + \kappa\Delta \quad \text{with} \quad \frac{d}{d\Delta}C_m(\Delta) = \kappa \geq 0.$$

On the down-wedge ($\Delta < 0$), the marginal cost is strictly negative and strictly decreasing as the lie moves further downward:

$$C_m(\Delta) = -\eta + \kappa\Delta \leq -\eta < 0, \quad \frac{d}{d\Delta}C_m(\Delta) = \kappa \geq 0.$$

Hence $|C_m(\Delta)| \geq \eta$ for all $\Delta \neq 0$, and $|C_m(\Delta)|$ is (weakly) increasing in $|\Delta|$.

S6 (Tail-growth). The signaling cost depends on the gap $\Delta := m - \theta$ via a convex, wedge-symmetric map $C(\Delta)$ with $C(0) = 0$ and C strictly increasing in $|\Delta|$. If for every θ the feasible message set allows a sequence $\{m_n\}$ with $|m_n - \theta| \to \infty$, then $C(\Delta) \to \infty$ as $|\Delta| \to \infty$, so sufficiently large lies are eventually prohibitive. When the feasible message set implies a bounded gap envelope, the tail-growth condition may not trigger separation and a finite-bar dominance test is required instead; see Online Appendix C for details.

S7: Type derivative: Because C depends on (m, θ) only through $\Delta = m - \theta$, we have $C_\theta = -C_m$ for $\Delta \neq 0$. No sign restriction on C_θ is imposed, and all comparative-statics that follow refer directly to C_m on the relevant wedge.

Assumptions (S1)–(S7) say that exaggeration is costly and the cost grows convexly with the size of the lie, with a discrete jump at truth-telling: even a tiny lie creates non-trivial enforcement exposure. Crucially, these properties are not assumed as primitives but derived from the Stage 2 enforcement technology (Section 4), where η and κ inherit their values from the ratification threshold r and judicial scrutiny σ.

This distinguishes the model from standard cheap-talk or costly-signaling setups (cf. Kartik, 2009): the cost function has institutional content. Its parameters encode the difficulty of buying peace through ratification (r-sensitive) and the probability that courts refuse to credit the outcome (σ-sensitive). When the carry-forward baseline links r to π, tightening the vote simultaneously steepens the cost surface—an interaction no one-stage model can generate. Two further consequences follow. First, enforcement costs are only partially deadweight: some fraction represents transfers, so the welfare cost of pooling is governed by the deadweight parameter ω. Second, changing the screening rule simultaneously changes the cost surface—impossible when signaling costs are exogenous.

4 Stage 2: Litigation, settlement, and ratification technology

This section describes the enforcement mechanism that operates after Stage-1 approval and derives the reduced-form enforcement cost $L(g, r, \sigma)$ used in (S1)–(S2).

Define the *misrepresentation gap* as $g := m - \theta$ (equivalently, $g := m^{\text{req}}(u) - \theta$ for the minimal persuasive lie). Let $r \in (0, 1]$ denote the *ratification/settlement vote threshold*: the fraction of non-conflicted shareholders whose approval is needed to extinguish claims. In the **carry-forward baseline**, $r = \pi$ (the same quota as the Stage-1 approval rule); more generally, $r = r(\pi)$ with $r'(\pi) \geq 0$.

After approval and extraction, minority shareholders may file suit. The mechanism works as follows. First, the *probability of suit* $p(g, r)$ is increasing in $|g|$: a larger gap between the controller's disclosure and the true project quality constitutes stronger prima facie evidence of misrepresentation, which encourages plaintiff attorneys and reduces the controller's likelihood of early dismissal. Second, the controller can attempt to extinguish claims by buying shareholder support for a *ratification vote* at threshold r (analogous to familiar cleansing logic for informed, disinterested stockholder votes). Ratification requires the controller to compensate enough minority shareholders to assemble an r-supermajority—effectively a vote-buying cost that is increasing in r because more shareholders must be persuaded. Third, courts may decline to credit a ratification vote as cleansing when the vote is not fully informed and uncoerced; in the model, this is captured by a deterministic credit discount $\sigma \in [0, 1)$, so only fraction $(1 - \sigma)$ of purchased votes is credited. Larger σ therefore raises the expected enforcement burden and can make cleansing infeasible when $r > 1 - \sigma$.

The distinctive feature of Stage 2 is that the controller actively buys peace through ratification, rather than passively waiting to be sued. Three primitives interact multiplicatively: a larger gap g raises the probability of suit and makes ratification harder; a tougher threshold r forces the controller deeper into the shareholder resistance distribution; and stricter scrutiny σ means purchased votes are more likely discounted. All three reinforce each other—large gaps are costliest when the threshold is high and scrutiny is strict—producing the convexity and supermodularity formalized below.

Solving the Stage-2 subgame backward yields a reduced-form expected enforcement cost $L(g, r, \sigma)$. Its properties follow from the mechanism above and inherit the supermodularity structure standard in complementarity economics (see Amir, 2005, for a survey):

Assumption 2 (Properties of the enforcement cost $L(g, r, \sigma)$)**.** *The enforcement cost function* $L : \mathbb{R} \times (0, 1] \times [0, 1) \to \mathbb{R}_+$ *satisfies:*

- (L1) $L(0, r, \sigma) = 0$, $L_g(g, r, \sigma) > 0$ *for* $g > 0$*, and* $L_{gg}(g, r, \sigma) \geq 0$.
  Why convex in g: *truthful disclosure* ($g = 0$) *carries no enforcement cost. A small misrepresentation creates a modest litigation exposure, but further exaggeration has an accelerating effect: the probability of suit rises, the expected damages grow, and the feasibility of settling via ratification falls—each of these channels reinforces the others, producing a convex cost surface.*
- (L2) $L_r(g, r, \sigma) \geq 0$ *for* $g > 0$*, and* $L_{gr}(g, r, \sigma) \geq 0$.
  Why $L_r \geq 0$: *a higher ratification threshold* r *forces the controller to buy more votes, each at increasing marginal cost, while simultaneously raising the probability that the court invalidates the ratification (since the assembled majority is more fragile). Both effects raise the expected enforcement cost for any given* $g > 0$*.*
  Why $L_{gr} \geq 0$ (supermodularity): *the marginal cost of an additional unit of misrepresentation is amplified when ratification is harder. Intuitively, a larger gap both strengthens the plaintiffs' legal position (making settlement more costly) and makes it harder to assemble a credible ratification vote (because more shareholders must be compensated at a higher per-vote price when* r *is large).*
- (L3) $L_\sigma(g, r, \sigma) \geq 0$ *for* $g > 0$*, and* $L_{g\sigma}(g, r, \sigma) \geq 0$.
  Why $L_\sigma \geq 0$: *a higher court-independence parameter* $\sigma \in [0, 1)$ *means paid votes are more likely to be disregarded, reducing the probability of successful ratification and raising the expected enforcement cost for any given misrepresentation.*
  Why $L_{g\sigma} \geq 0$ (supermodularity in g and σ): *when courts scrutinize vote independence more strictly, the marginal cost of an additional unit of misrepresentation rises—the controller must offer larger side payments to compensate for the higher probability of judicial invalidation,*

and these payments grow faster in g when σ is large.

As a robustness extension, court-credit risk can be made endogenous in the disclosure gap: $\sigma = \sigma(g)$ with $\sigma'(g) \geq 0$, capturing the intuition that courts are less likely to credit ratification votes when the disclosure gap is large. This extension amplifies the litigation channel and relaxes the conditions needed for LA-DWBR; details are in Online Appendix C. For tractability, all equilibrium characterizations and proofs in the main text keep σ exogenous.

The functional form $L(g, r, \sigma) = \eta(r, \sigma)|g| + \frac{\kappa(r,\sigma)}{2} g^2$ with $\eta_r \geq 0$, $\kappa_r \geq 0$, $\eta_\sigma \geq 0$, and $\kappa_\sigma \geq 0$ satisfies (L1)–(L3) and nests (S1)–(S2). The parameter $\eta(r, \sigma)$ captures the *fixed marginal exposure* from any non-zero misrepresentation—the per-unit cost of litigation risk at the origin of the gap—and increases in both r and σ because a higher ratification threshold or more stringent independence scrutiny makes it costlier to buy peace even for a small misrepresentation. The parameter $\kappa(r, \sigma)$ governs the *convexity* of the cost in the gap, reflecting the accelerating difficulty of defending or settling larger misrepresentations; $\kappa_r \geq 0$ and $\kappa_\sigma \geq 0$ because harder ratification or stricter scrutiny makes the settlement probability fall faster as the gap grows.

When extraction is endogenous ($x \in [0, B(\theta)]$), the reduced-form net cost is isomorphic to the fixed-extraction specification; the continuous formulation is in Online Appendix C.

We now derive the ratification cost from an explicit vote-buying game. Suppose each minority shareholder i has privately known ratification-resistance type $v_i \overset{\text{iid}}{\sim} G$ on $[\underline{v}, \bar{v}]$, where G admits a continuous density $g > 0$ on $(\underline{v}, \bar{v})$ and satisfies $G(\underline{v}) = 0$, $G(\bar{v}) = 1$. Let $R \in [0, 1]$ denote the per-share recovery at stake if litigation continues. If shareholder i votes to ratify, she forgoes continuation value and bears resistance cost; with transfer $t_i \geq 0$, utility from ratifying is

$$u_i^Y = t_i - (R + v_i),$$

while utility from rejecting is normalized to $u_i^N = 0$. Hence i votes to ratify iff

$$t_i \geq R + v_i.$$

The reservation price $R+v_i$ has two economically distinct components. The common component R compensates for the litigation value that every shareholder forgoes by ratifying—the expected per-share recovery if the lawsuit continued to judgment, which depends on the strength of the legal case and hence on the misrepresentation gap g. The idiosyncratic component v_i compensates for shareholder i's private resistance to settlement. The additive structure means that an increase in R shifts every shareholder's reservation price upward by the same amount—a parallel shift of the supply curve—while the heterogeneity in v_i generates the curve's upward slope.

The controller therefore buys support from the lowest-reservation-price voters first (lowest v_i), with minimum transfer schedule

$$t(v) = \max\{0,\ R + v\}.$$

The resistance type v_i captures idiosyncratic reluctance to ratify: passive index funds have low v_i; activist hedge funds have high v_i. The distribution G reflects shareholder-base composition (cf. Hamdani and Yafeh, 2013; Hu and Black, 2006). Crucially, the qualitative results below do not depend on the shape of G—all key derivative signs are distribution-free (Remark E.4.4).

For tractability, we adopt a deterministic court-credit discount: only a fraction $(1-\sigma)$ of purchased votes is credited as cleansing, with $\sigma \in [0, 1)$. If the controller buys fraction q of minority votes, the credited fraction is $(1 - \sigma)q$. To satisfy ratification threshold r, he must buy at least

$$q \;\geq\; \frac{r}{1-\sigma}.$$

Purchasing fraction q means buying all types $v \leq v^*$ where $G(v^*) = q$, i.e. $v^* = G^{-1}(q)$. Thus ratification is feasible only on

$$0 < r \leq 1 - \sigma.$$

When $r > 1 - \sigma$, no amount of vote buying can produce enough credited votes.

Lemma 1 (Ratification cost – general resistance distribution)**.** *Under the vote-buying model above with $R \in [0, 1]$, $\sigma \in [0, 1)$, threshold $r \in (0, 1]$, and resistance distribution G with continuous positive density g on $(\underline{v}, \bar{v})$. Define the vote-buying cutoff*

$$v^*(r, \sigma) \;:=\; G^{-1}\Big(\frac{r}{1-\sigma}\Big),$$

and the kink threshold

$$r_0(R, \sigma) \;:=\; (1-\sigma)\, G(-R),$$

with the convention that $r_0 = 0$ if $-R \leq \underline{v}$.

The controller's minimum ratification expenditure is:

$$C(R, r, \sigma) \;=\; \begin{cases} \infty, & r > 1 - \sigma, \\ 0, & 0 < r \leq r_0(R, \sigma), \\ \displaystyle\int_{\max\{\underline{v},\, -R\}}^{v^*(r,\sigma)} (R + v)\, g(v)\, dv, & r_0(R, \sigma) < r \leq 1 - \sigma. \end{cases} \tag{1}$$

On the active interior region $r_0(R, \sigma) < r < 1 - \sigma$ where $R + v^* > 0$, the following sign conditions hold:

(a) $C_R \;=\; G(v^*) - G\big(\max\{\underline{v}, -R\}\big) \;>\; 0,$ (b) $C_{RR} \;=\; g(-R)\, \mathbf{1}\{-R > \underline{v}\} \;\geq\; 0,$

(c) $C_r \;=\; \dfrac{R + v^*}{1-\sigma} \;>\; 0,$ (d) $C_\sigma \;=\; \dfrac{r\,(R + v^*)}{(1-\sigma)^2} \;>\; 0,$

(e) $C_{Rr} \;=\; \dfrac{1}{1-\sigma} \;>\; 0,$ (f) $C_{R\sigma} \;=\; \dfrac{r}{(1-\sigma)^2} \;>\; 0.$

In particular, $C_{RR} > 0$ whenever $-R \in (\underline{v}, \bar{v})$, i.e. whenever some voters have negative reservation *values. At the kink $r = r_0(R, \sigma)$, right-derivatives in r are zero and all inequalities hold weakly.*

Proof. Proof in Appendix E.4, §E.4.1. □

The controller faces an upward-sloping supply curve of ratification votes: more votes (higher r) or greater judicial discounting (higher σ) force him deeper into the resistance distribution. The key takeaway is that all comparative-statics signs are distribution-free: the density $g(v^*)$ cancels in every Leibniz evaluation at the pivotal voter, so the marginal cost of a tougher threshold is governed by the pivotal voter's reservation price $R + v^*$, not the density of voters near the margin (Remark E.4.4).

Below the kink ($r \le r_0$), the controller extinguishes claims at zero cost—enforcement is toothless. The disciplinary force of Stage 2 activates only when r pushes past the kink into the active region. This is the core argument for coupling: under $r(\pi) = \pi$, raising the approval bar simultaneously pushes enforcement from the zero-cost into the active region, so the two margins move together.

Remark (Distribution-free signs). The sign conditions $C_r > 0$, $C_\sigma > 0$, $C_{Rr} > 0$, and $C_{R\sigma} > 0$ hold for any continuous resistance distribution G with positive density. The density $g(v^*)$ cancels in every Leibniz evaluation because the marginal cost of a tougher threshold is determined by the pivotal voter's reservation price $R + v^*$, not the density of voters near the margin. The one exception is $C_{RR} = g(-R)\,\mathbf{1}\{-R > \underline{v}\}$, which depends on the density at the kink rather than at the pivotal voter. See Appendix D for the full economic argument.

Under uniform resistance $v_i \sim \text{Unif}[-1, 1]$, the ratification cost admits a closed-form quadratic expression with all derivative signs confirmed; see Appendix D for the statement and proof.

Define the *recovery at stake* as $R(g, x) = \ell(g) \cdot x$, where $\ell : \mathbb{R}_+ \to \mathbb{R}_+$ is increasing and convex in $|g|$ (litigation exposure per unit extracted). The reduced-form enforcement cost is then

$$L(g, r, \sigma) \;:=\; p(g, r) \cdot C\big(R(g, x^*),\; r,\; \sigma\big),$$

where $p(g, r)$ is the probability of suit and x^* is the optimal extraction. On the feasible cleansing region $r \le 1 - \sigma$, C is increasing-convex in R and increasing in r, σ on the active branch. Proposition 1 gives explicit sufficient conditions under which the composite $L(g, r, \sigma)$ inherits (L1)–(L3) on the up-wedge $g \ge 0$. For $r > 1 - \sigma$, cleansing is infeasible and the effective enforcement burden is unbounded in this microfoundation.

Remark (Nesting). The parametric form $\eta(r, \sigma)|g| + \frac{\kappa(r,\sigma)}{2} g^2$ is a second-order Taylor expansion of $L(g, r, \sigma)$ around $g = 0$. All propositions stated under Assumption 2 therefore remain valid under the ratification-vote microfoundation; the microfoundation provides additional closed-form structure but does not alter the general results.

The reduced-form $L(g, r, \sigma)$ treats the gap as the sufficient statistic for enforcement costs, integrating over litigation paths—a standard approach (Bebchuk, 1984).

Proposition 1 (Ratification microfoundation implies (L1)–(L3) on the feasible region)**.** *Fix* $\sigma \in [0,1)$ *and define* $d := 1-\sigma$. *On the up-wedge* $g \geq 0$, *consider*

$$L(g,r,\sigma) \;:=\; p(g,r)\,C\big(R(g),r,\sigma\big), \qquad 0 < r \leq d,$$

where $C(R,r,\sigma)$ *is the ratification cost from Lemma 1 under a general resistance distribution* G *with continuous positive density* g *on* $(\underline{v},\overline{v})$. *Assume:*

(a) $p \in C^2$, $p(0,r)=0$, *and* $p_g, p_{gg}, p_r, p_{gr} \geq 0$ *on* $g \geq 0$, $0 < r \leq d$.
(b) $R \in C^2$, $R(0)=0$, *and* $R_g, R_{gg} \geq 0$ *on* $g \geq 0$.
(c) Either (i) $r > r_0(R(g),\sigma)$ *(active branch of* C*), or (ii) one-sided derivatives at the kink* $r = r_0$ *are used, giving weak inequalities.*

Then on $g \geq 0$, $0 < r \leq d$, L *satisfies* (L1)–(L3)*:*

$$L(0,r,\sigma)=0, \quad L_g \geq 0, \quad L_{gg} \geq 0, \quad L_r \geq 0, \quad L_{gr} \geq 0, \quad L_\sigma \geq 0, \quad L_{g\sigma} \geq 0.$$

Strict inequalities hold on the active interior whenever at least one corresponding primitive derivative $(p_g, p_{gg}, p_r, p_{gr}, R_g$, *or* $R_{gg})$ *is strictly positive.*

Proof. Proof in Appendix E.4, §E.4.2. □

An alternative settlement-bargaining microfoundation (cf. Bebchuk, 1984) confirms (L1)–(L3) under standard monotonicity conditions; see Online Appendix C.

4.0.1 Controlling Shareholder's Utility Function

Given the primitives in Sections 3.2.1 and 3.2.3, and the voting cut-off in (V2), define the rule-specific approval indicator

$$\mathbf{1}\{\text{approve}(m)\} \;:=\; \mathbf{1}\Big\{(1-\alpha)\,\mu(m) \;-\; \beta^{\text{Div}}(m) \;\geq\; \kappa^{\text{rule}}(\lambda)\Big\} \;=\; \mathbf{1}\{\mu(m) \;\geq\; u\}\,.$$

where u is the belief cut-off from (V3). The controlling shareholder chooses a message anticipating minority voting and the expected Stage-2 enforcement cost:

$$\max_{m}\; U_{CS}(m,\theta) \;=\; \mathbf{1}\{\text{approve}(m)\}\cdot\big(\alpha\,\theta + B(\theta)\big) \;-\; C(m,\theta), \tag{C1}$$

where $C(m,\theta) \equiv L(g,r,\sigma)$ is the reduced-form enforcement burden from Stage-2 litigation (Section 4), with $g = m-\theta$ the misrepresentation gap, r the ratification threshold, and σ the court-independence parameter. When r and σ are fixed, the Stage-1 problem is identical to the standard signaling model; the novel interaction arises under the carry-forward baseline $r = \pi$.

4.0.2 Controlling Shareholder's Participation Constraint

Absent additional structure, Assumption P3 together with the payoff criterion in (C1) implies that a controller has little incentive to initiate projects with $\theta \leq 0$. To ensure well-posed participation at the voting cut-off, we adopt the affine specification $B(\theta) = b_0 + b_1\theta$ with $b_1 \in [0, 1 - \alpha)$ and choose b_0 large enough that the controller finds it (weakly) profitable to proceed at the bar.[4] This is encoded in the following assumption.

Assumption 3 (Controller participation at the rule cut-off (CPC) and interior scope)**.** *Maintain Assumptions P1–P3, S1–S4, and M1. In the baseline, take $B'(\theta) \in (0, 1 - \alpha)$ and $B(\theta) = b_0 + b_1\theta$ with $b_1 \in [0, 1-\alpha)$. For any voting environment that induces a belief cut-off u via* (V2)-(V3)*, require the controller's participation constraint at the cut-off u to hold:*

$$\alpha\, u + B(u) \;\geq\; 0. \tag{CPC}$$

Under $B(\theta) = b_0 + b_1\theta$ this is equivalent to

$$(\alpha + b_1)\, u + b_0 \;\geq\; 0 \qquad \Longleftrightarrow \qquad u \;\geq\; u_{\text{CPC}} \;:=\; -\frac{b_0}{\alpha + b_1}.$$

Because $\alpha + b_1 > 0$, (CPC) imposes a lower bound u_{CPC} on admissible cut-offs. In the baseline, b_0 is chosen so that $u \geq u_{\text{CPC}}$ along every interior corridor of cut-offs considered, ensuring the controller participates at the bar. Notably, under $u_{\text{MoM}} = 0$, CPC reduces to $b_0 \geq 0$; the pure perk can alone ensure participation at the MoM bar.

5 Stage 1 equilibrium given Stage 2 enforcement costs

5.1 Minority Shareholders

5.1.1 Minority Shareholder Voting and Posterior Beliefs

As discussed in Section 3.2.1, a large but finite number N of symmetric minority shareholders collectively hold the residual $1 - \alpha$ of the project's valuation, θ. To simplify notation, I define the variables s and λ:

$$s \;=\; \frac{1-\alpha}{N}, \qquad \lambda \;:=\; \frac{s}{1-\alpha} \;=\; \frac{1}{N} \;\in\; (0, 1).$$

Hence each minority shareholder owns the fraction λ of the residual stake and, correspondingly, is exposed to a λ-proportion of any private benefits extracted by the controlling shareholder. We assume $\sum_{n=1}^{N} \lambda\, D(\theta) = D(\theta)$, i.e., only the *diversionary* (zero-sum) component is borne pro-rata by minorities; the pure perk b_0 does not reduce minorities' payoff. Except where noted, we

[4] Henceforth, we refer to the "bar" as the belief cut-off, u defined in V3.

will henceforth treat the case of a single representative minority shareholder, without loss of generality.

5.1.2 Beliefs and Updating

The equilibrium throughout is a Perfect Sequential Equilibrium in which the controller (sender) is fully strategic and minority shareholders (receivers) are atomistic. Each shareholder observes the disclosed message m, forms belief $\mu(m)$ using the effective posterior defined below, and votes according to the aggregation rule (V2). Shareholders do not reverse-engineer the controller's equilibrium strategy to update further. This modeling choice is justified by the partial-attention microfoundation developed immediately below: empirically, most minority shares are not actively voted, and the resulting effective posterior can sustain the pooling equilibrium even under Bayesian updating by attentive shareholders. Appendix E.4.5 provides a detailed discussion.

For each minority shareholder, a *belief system* specifies, for each publicly observed message $m \in \mathbb{R}$, a posterior distribution $\rho(\cdot \mid m) \in \Delta([\underline{\theta}_0, \overline{\theta}_0])$ over types. On the equilibrium path, beliefs are obtained by Bayes' rule from the prior F and the sender's message strategy $m(\cdot)$;[5] off the equilibrium path we adopt the Grossman–Perry *credible-beliefs* refinement (Appendix A.4): upon observing an unexpected message, minorities place probability 1 on the set of types that (weakly) benefit most from that deviation (Grossman and Perry, 1986).

We summarize the fully attentive Bayesian posteriors by their expectations:

$$\bar{\mu}(m) := \mathrm{E}[\theta \mid m], \qquad \beta(m) := \mathrm{E}[B(\theta) \mid m], \qquad \beta^{\mathrm{Div}}(m) := \mathrm{E}[D(\theta) \mid m].$$

The controller's payoff uses $\beta(m)$ (full B); the voting test for minorities uses $\beta^{\mathrm{Div}}(m)$ (diversion only).

Not all shareholders are attentive. Let $\tau \in (0, 1]$ denote the fraction of minority shares (equivalently, voters) that process the controller's message m.[6] The remaining $1 - \tau$ are inattentive, abstain, or vote mechanically using the prior mean $\mu_0 := \mathrm{E}[\theta]$. The *effective posterior mean* that enters the voting aggregator is

$$\mu(m) \;:=\; \tau\,\bar{\mu}(m) + (1 - \tau)\,\mu_0.$$

Throughout the paper, $\mu(m)$ denotes this effective voting statistic, not the fully attentive Bayesian posterior $\bar{\mu}(m)$. The mixture implies $\mu'(m) = \tau\,\bar{\mu}'(m)$, so responsiveness is scaled down when attention is low. The benchmark is nested at $\tau = 1$ (so $\mu = \bar{\mu}$).

[5] See Section 5.2.5 for strategy notation.

[6] Broadridge ProxyPulse reports that retail investors vote only a minority of the shares they own (roughly 28–30%, including 28% in the 2025 update), far below institutional voting rates; see https://www.broadridge.com/_assets/pdf/2025proxypulse-updated.pdf. See also Brav et al. (2022) and Lauterbach and Mugerman (2020).

We assume that:

$$\mu(m) \text{ is } C^1 \text{ with } \mu'(m) > 0 \text{ for all } m; \text{ accordingly, } \mu^{-1} \text{ exists and is } C^1 \text{ with } (\mu^{-1})'(u) = 1/\mu'(\mu^{-1}(u)) > 0. \tag{M1}$$

We further impose the following tail-limit regularity on beliefs: as messages become arbitrarily small or large, the posterior mean converges to the endpoints of the valuation support, i.e.,

$$\lim_{m \to -\infty} \mu(m) = \underline{\theta}_0, \qquad \lim_{m \to +\infty} \mu(m) = \overline{\theta}_0. \tag{M1$'$}$$

5.1.3 Endogenous participation and turnout (robustness summary)

The baseline keeps belief responsiveness in reduced form. A turnout microfoundation can be added without changing any core results: let attentive participation be $\tau(m) \in (0,1]$, with attentive voters using $\bar{\mu}(m) = \mathrm{E}[\theta \mid m]$ and inattentive voters using the prior $\mu_0 = \mathrm{E}[\theta]$. The effective voting belief is

$$\mu(m) = \tau(m)\,\bar{\mu}(m) + \big(1 - \tau(m)\big)\mu_0.$$

When votes are not expected to be close, participation is low and beliefs are sluggish; near the bar, pivotality raises participation and responsiveness. This interpretation supports the same comparative statics and maps the local responsiveness parameter in Appendix C to attention-weighted slope near the boundary (cf. Meirowitz and Pi, 2022, for a related voting-trading dilemma that generates similar sluggishness). Detailed fixed-point turnout derivations are reported in Appendix B and the Online Appendix.

5.1.4 Expected Utility from Approving (Interim Stage)

Following Levit et al. (2024) and consistent with evidence on institutional investors' limited engagement in concentrated-ownership settings (Hamdani and Yafeh, 2013), we endow each minority investor i with a privately known bias b_i that perturbs her cut-off for approving the proposal. The bias is normalised as

$$b_i \;\sim\; \text{Uniform}\,[-1,\,1], \qquad \mathbb{E}[b_i] = 0, \tag{M2}$$

independent of θ and distributed across investors. Using the posterior summaries from Section 5.1.2, investor i's interim utility upon approving is

$$\mathbb{E}[u_i \mid m] = \lambda\Big[(1-\alpha)\,\mu(m) - \beta^{\text{Div}}(m)\Big] + b_i, \tag{M3}$$

with $G(b) = (b+1)/2$ for $b \sim U[-1,1]$.

5.1.5 Yes/No rule and Pivotal bias

We define

$$t_\lambda(m) := -\lambda\big[(1-\alpha)\mu(m) - \beta^{\text{Div}}(m)\big].$$

Thus, a minority shareholder i votes "Yes" if and only if $b_i \geq t_\lambda(m)$. Accordingly, the fraction of "Yes" votes is

$$1 - G\big(t_\lambda(m)\big) = \frac{1 - t_\lambda(m)}{2}, \tag{M4}$$

given the distribution of $b_i \sim U[-1, 1]$, for $t_\lambda(m) \in [-1, 1]$ ($t_\lambda(m)$ saturates at 0 or 1 outside).

5.2 Voting Rules

5.2.1 Definition

We consider two types of Voting Rules.

1. **Majority of the Minority (MoM).** A proposal is implemented if and only if strictly more than 0.5 of the $(1-\alpha)$ minority shares vote "Yes."

2. **Super-majority (SM).** Fix a quota $S \in (0.5, 1)$. Implementation requires at least a fraction S of *all* shares in favour. Equivalently, the fraction of minority votes needed is

$$\pi = \frac{S - \alpha}{1 - \alpha}. \tag{V1}$$

   An SM rule is defined to be *lenient* when $\pi < 0.5$ and *stringent* when $\pi > 0.5$.

With these primitives fixed, all subsequent analysis derives equilibrium signalling and voting outcomes and compares welfare across the two rule families.

5.2.2 Cut-offs under MoM and super-majority S

For the MoM rule, the cutoff point is defined by the minority shareholder who is indifferent between voting "Yes" and "No" for a given proposal, conditional on the message m sent by the CS. As such, given expressions M4 and V1, under the MoM rule, shareholders vote "Yes" iff:

$$\frac{1 - t_\lambda(m)}{2} \geq \tfrac{1}{2} \implies t_\lambda(m) \leq 0 \implies (1-\alpha)\,\mu(m) - \beta^{\text{Div}}(m) \geq 0.$$

Similarly, under the SM rule, let $\pi = (S - \alpha)/(1 - \alpha)$. Requiring $(1 - t_\lambda(m))/2 \geq \pi$ as per expressions M4 and V1 gives the condition where shareholders vote "Yes" iff:

$$t_\lambda(m) \le 1 - 2\pi \implies (1-\alpha)\,\mu(m) - \beta^{\mathrm{Div}}(m) \;\ge\; -\frac{1-2\pi}{\lambda}.$$

We can then define the rule-specific cutoff

$$\underline{\kappa}^{(\lambda)}_{\mathrm{rule}} := -\frac{k_{\mathrm{rule}}}{\lambda}, \quad k_{\mathrm{MoM}} = 0, \quad k_{\mathrm{SM}} = 1 - 2\pi,$$

so that in both cases minorities approve iff:

$$(1-\alpha)\,\mu(m) - \beta^{\mathrm{Div}}(m) \;\ge\; \underline{\kappa}^{(\lambda)}_{\mathrm{rule}}. \tag{V2}$$

5.2.3 Belief Cut-offs u

Fix a voting rule and exposure λ. Let $\kappa^{(\lambda)}_{\mathrm{rule}}$ be the constant in (V2). Define the minimal persuasive message which would pass the proposal given a voting rule:

$$m_{\mathrm{req}} := \inf\left\{\, m \in \mathbb{R} \;:\; (1-\alpha)\mu(m) - \beta^{\mathrm{Div}}(m) \;\ge\; \kappa^{(\lambda)}_{\mathrm{rule}} \,\right\}, \qquad u := \mu(m_{\mathrm{req}}). \tag{V3}$$

By (M1), μ is strictly increasing and continuous, so m_{req} exists and is unique with $m_{\mathrm{req}} = \mu^{-1}(u)$. Given that $\mu(m) = \mathbb{E}[\theta \mid m]$, u is the conditional expected valuation of the pivotal shareholder who induces the proposal to pass. We refer to u as the belief cut-off/threshold for passage.

The voting outcome is governed by the primitive belief function μ satisfying (M1), not by the equilibrium Bayesian posterior. In the atomistic-voter framework of Levit et al. (2024), each shareholder observes m, forms $\mu(m)$, and votes according to (V2) without reverse-engineering the controller's strategy. See Appendix D for a detailed discussion.

5.2.4 Regime Orientation

When diversion is modest, $0 < B'(\theta) < 1 - \alpha$ (*payoff-dominant, talk-up*), minorities' net payoff $(1-\alpha)\theta - D(\theta)$ rises with true quality, so they welcome higher perceived value. The approval test is $\mu(m) \ge u$, any pooling (if present) lies *below* u, and persuasion is *upward.* Equivalently, with $\Phi(m) := (1-\alpha)\mu(m) - \beta^{\mathrm{Div}}(m)$, Φ is strictly increasing in μ, yielding a unique interior u.

When diversion is steep, $B'(\theta) > 1 - \alpha$ (*diversion–dominant*), the approval test flips to $\mu(m) \le u$ and persuasion is downward. This extension is deferred from the main text.

5.2.5 Interior cut-offs and strategy notation

Standing assumption (interior cut-off): Unless stated otherwise, $u \in (\underline{\theta}_0, \overline{\theta}_0)$.[7] The degenerate cases $u \notin (\theta_0, \bar{\theta}_0)$ (all-Yes/all-No) are collected in Corollary 4 (Appendix A).

Strategy notation: The controller's type-contingent message strategy is a function $m_u : [\underline{\theta}_0, \overline{\theta}_0] \to \mathbb{R}$, written $m_u(\theta)$; when the u-dependence is not important, we write $m(\theta)$.

5.2.6 Closed-Form Cut-Offs with Affine Private Benefits

Assume $B(\theta) = b_0 + b_1\theta$ with $b_1 \in (0,\ 1-\alpha)$. Write the diversionary part as $D(\theta) = b_1\theta$ (so b_0 is a pure perk). Given any message m,

$$\beta^{\mathrm{Div}}(m) \;=\; \mathrm{E}\big[D(\theta) \mid m\big] \;=\; b_1\,\mu(m).$$

Substituting into the common voting inequality (V2),

$$(1-\alpha)\,\mu(m) \;-\; \beta^{\mathrm{Div}}(m) \;=\; (1-\alpha-b_1)\,\mu(m) \;\geq\; \underline{\kappa}^{(\lambda)}_{\mathrm{rule}}.$$

Because $1-\alpha-b_1 > 0$, this is equivalent to the belief threshold

$$\mu(m) \;\geq\; u_{\mathrm{rule}}(\lambda) \;:=\; \frac{\underline{\kappa}^{(\lambda)}_{\mathrm{rule}}}{1-\alpha-b_1} \tag{V4}$$

for each voting rule.

Recalling $\underline{\kappa}^{(\lambda)}_{\mathrm{MoM}} = 0$ and $\underline{\kappa}^{(\lambda)}_{\mathrm{SM}} = -(1-2\pi)/\lambda$, (V4) gives

$$u_{\mathrm{MoM}} = 0, \qquad u_{\mathrm{SM}}(\pi,\lambda) \;=\; -\frac{1-2\pi}{\lambda(1-\alpha-b_1)}.$$

Thus, the pure perk b_0 does not enter the minorities' bar: it affects only the controller's prize and the CPC constraint below.

Keeping (α, b_1) fixed,

$$\frac{\partial u_{\mathrm{SM}}}{\partial \pi} = \frac{2/\lambda}{1-\alpha-b_1} > 0.$$

Hence stricter quotas (higher π) tighten the belief cut-offs u required to pass a given proposal in proportion to $(1-\alpha-b_1)^{-1}$. In a separating equilibrium (so that $\mu(m) = \theta$), the belief thresholds above translate directly into type thresholds, $\theta \geq u_{\mathrm{rule}}(\lambda)$. In a pooling equilibrium, minorities compare the pooled posterior mean with $u_{\mathrm{rule}}(\lambda)$ from V4.

[7] The standing interior-cutoff assumption ensures that the analysis operates away from degenerate atomistic limits. As $N \to \infty$, lenient supermajority rules approach universal approval and stringent rules approach universal rejection, while MoM is neutral; these are boundary artifacts of the $\lambda = 1/N$ scaling that do not affect comparative statics on the interior of the policy domain. The paper's results are stated for fixed finite N with interior u.

5.3 Equilibrium structure at a fixed bar

We now characterize Stage-1 behavior taking Stage-2 enforcement technology as given. The key object is the lying-profit function, which values approval net of expected enforcement costs. We begin with lemmas that apply to *any* voting rule in the baseline talk-up regime.

5.3.1 Lying Profit Function

We reserve $u \in (\underline{\theta}_0, \overline{\theta}_0)$ for the rule-specific belief cut-off from (V3) and θ for types. First, we redefine the "minimal persuasive message" in V3 as the "cheapest effective lie" $m_{\text{req}}(u)$–the lowest-cost message required to induce approvals by the pivotal minority shareholders given their belief thresholds.

Second, define the CS's net gain from lying as a "lying-profit function", $\Lambda(\theta; u)$, which is a function both of the project's valuation θ and the cutoff u. A "lying-profit function" determines the payoff to the CS from misrepresenting the true valuation of the project, θ, such that minority shareholders will accept the proposal in equilibrium. We define $\Lambda(\theta; u)$ as:

$$\Lambda(\theta; u) \ := \ \alpha\,\theta + B(\theta) - C\big(m_{\text{req}}(u), \theta\big) \ = \ \alpha\,\theta + B(\theta) - L\big(m_{\text{req}}(u) - \theta,\ r, \sigma\big),$$

where the second equality uses $C(m, \theta) \equiv L(g, r, \sigma)$ with $g = m_{\text{req}}(u) - \theta$. When $\Lambda(\theta; u) > 0$ for types just below u, at least some such types strictly prefer to shade up to $m_{\text{req}}(u)$ to secure approval. When r varies with π (carry-forward baseline), Λ acquires an additional dependence on r through the enforcement cost:

$$\Lambda(\theta; u, r) \ := \ \alpha\,\theta + B(\theta) - L\big(m_{\text{req}}(u) - \theta,\ r, \sigma\big).$$

The lying-profit function asks whether a controller with true quality θ should exaggerate to the bar. High-quality types (close to u) face small gaps and low enforcement costs, so they pool. Low-quality types face large gaps and convex costs, so they separate truthfully and get rejected. The boundary type θ^* is exactly indifferent. Everything flows from how θ^* responds to the policy parameters: when the ratification threshold is coupled to the approval threshold, the enforcement cost surface steepens, pushing θ^* upward and shrinking the pool. Conditional on exaggerating, every pooler sends exactly $m_{\text{req}}(u)$: any higher message creates a strictly larger gap with zero additional benefit.

5.3.2 General Existence and Uniqueness

As noted by Burkart et al. (2024), signaling problems often induce multiple equilibria. We impose GP-credible off-path beliefs for uniqueness; supporting lemmas are in Appendix A.

Proposition 2 (Perfect-Sequential Equilibrium for a fixed belief cut-off)**.** *Fix any rule that induces* *the belief cut-off* $u \in (\underline{\theta}_0, \overline{\theta}_0]$. *Define* $m_{\mathrm{req}}(u) := \mu^{-1}(u)$ *and*

$$\Lambda(\theta; u) \ := \ \alpha\theta + B(\theta) - C\big(m_{\mathrm{req}}(u), \theta\big).$$

Let $\theta^*(u) := \sup\{\theta \in [\underline{\theta}_0, u) : \Lambda(\theta; u) < 0\}$. *Minorities implement iff* $\mu(m) \geq u$ *(equivalently, iff* $(1-\alpha)\mu(m) - \beta^{\mathrm{Div}}(m) \geq \underline{\kappa}^{(\lambda)}_{rule}$ *by* (V2)*). Assume GP-credible beliefs off path. Throughout this proposition assume* $0 < B'(\theta) < 1 - \alpha$ (P3).

(i) *Existence. There exists a PSE with sender message strategy*

$$m(\theta) = \begin{cases} \theta, & \theta < \theta^*(u) \quad \textit{(truth; rejected)}, \\ m_{\mathrm{req}}(u), & \theta \in [\theta^*(u), u) \quad \textit{(cheapest lie; approved)}, \\ \theta, & \theta \geq u \quad \textit{(truth; approved)}. \end{cases}$$

(ii) *Allocation uniqueness for a given* u*. For the same cut-off* u*, every PSE yields the same implement/reject allocation: the project is implemented if and only if* $\theta \geq \theta^*(u)$.

(iii) *Pooling vs. separation (endpoint test). Write*

$$\Psi_\ell(u) \ := \ \Lambda(\underline{\theta}_0; u), \qquad \Psi_h(u) \ := \ \lim_{\theta \uparrow u} \Lambda(\theta; u) \ = \ \alpha u + B(u) - C\big(m^{\mathrm{req}}(u), u\big).$$

Since $\theta \mapsto \Lambda(\theta; u)$ *is continuous and strictly increasing on* $[\underline{\theta}_0, u]$*, exactly one of the following obtains:*

(A) Partial pooling: $\Psi_\ell(u) < 0 < \Psi_h(u) \implies \exists!\ \theta^*(u) \in (\underline{\theta}_0, u),\ F(u) - F(\theta^*(u)) > 0.$

(B) Full separation: $\Psi_h(u) \leq 0 \implies \theta^*(u) = u$ *and the pooling set is empty.*

(C) Full pooling below u*:* $\Psi_\ell(u) \geq 0 \implies \theta^*(u) = \underline{\theta}_0$ *and all* $\theta < u$ *pool at* $m_{\mathrm{req}}(u)$.

Proof sketch in main text; full proof in Appendix A.1. The fixed-u allocation follows from strict single-crossing of the lying-profit function, endpoint tests at $\underline{\theta}_0$ and u, and GP-credible beliefs. A complete type-by-type deviation argument is reported in Appendix A.1. □

The equilibrium has a three-region structure. Bottom types ($\theta < \theta^*$) tell the truth and get rejected. Middle types ($\theta \in [\theta^*, u)$) pool at $m_{\mathrm{req}}(u)$—they exaggerate just enough to clear the bar. Top types ($\theta \geq u$) tell the truth and get approved. Single-crossing ensures this partition is contiguous: the lying-profit function is increasing in true quality, so the set of exaggerators is always a connected interval.

The pool has a skewed cost structure. Every pooler sends $m_{\mathrm{req}}(u)$, so the gap $g(\theta) = m_{\mathrm{req}}(u) - \theta$ is largest at the boundary θ^* and approaches zero near u. By convexity of L, bottom-of-pool types are disproportionately expensive. When coupling raises θ^*, it trims the pool from the expensive end, making bottom-trimming especially effective at reducing enforcement waste—the mass effect formalized in Section 6.

Figure 1: Equilibrium message policy and per-type enforcement cost

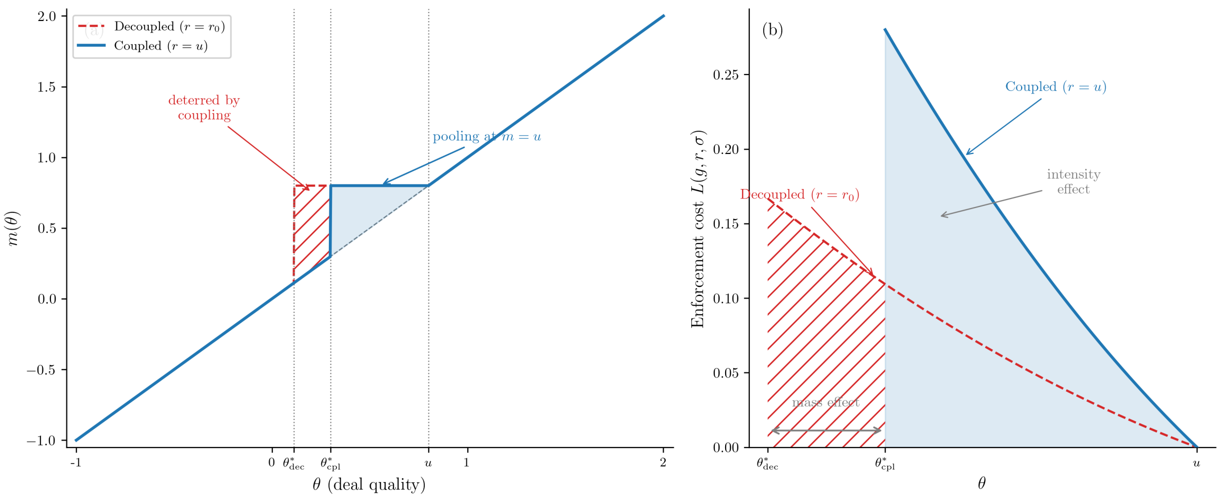

Panel (a): the equilibrium disclosure function $m(\theta)$ under coupled (solid blue, $r = u$) and decoupled (dashed red, $r = r_0$) regimes at bar $u = 0.8$. Types below θ^* separate truthfully and are rejected; types in $[\theta^*, u)$ pool at $m = u$; types above u separate truthfully and are approved. Coupling raises θ^* (from θ^*_{dec} to θ^*_{cpl}), deterring the hatched types who would have pooled under the lax regime. Panel (b): per-type enforcement cost $L(g, r, \sigma)$ within the pooling region, where $g = u - \theta$. The coupled curve (solid) is steeper—each unit of gap is costlier under tough ratification (intensity effect)—but starts further right because the most expensive types are deterred (mass effect). Parameters: $\alpha = 0.4$, $b_0 = 0.1$, $b_1 = 0.2$, $\sigma = 0.2$, $r_0 = 0.5$.

5.4 Main Result: Carry-Forward Ratification Reverses Pooling Comparative Statics

We now state the paper's central result. Under the carry-forward baseline $r(\pi) = \pi$, the pooling derivative $dP/d\pi$ acquires a *ratification-channel* term absent from any one-stage model. When this term is large enough, it reverses the benchmark prediction that tighter votes increase misrepresentation.

Proposition 3 (Litigation-adjusted pooling comparative statics). *Maintain* (P1)–(P3), (S1)–(S7), *Assumption 2, and* (M1) *in the payoff-dominant (up-talk) regime* $0 < B'(\theta) < 1 - \alpha$. *Let* $P(\pi) := \int_{\theta^*(u(\pi),r(\pi))}^{u(\pi)} f(\theta)\, d\theta$ *denote the pooling mass under approval rule* π, *where* $u(\pi)$ *is the induced belief cut-off and* $r(\pi)$ *the ratification threshold with* $r'(\pi) \geq 0$. *Define* $g^*(\pi) := m^{\text{req}}\big(u(\pi)\big) - \theta^*$ *as the equilibrium misrepresentation gap at the boundary. If one specializes to the ratification-vote microfoundation in Lemma 1, statements below are evaluated on the feasible corridor* $r(\pi) \leq 1 - \sigma$.

(i) *The pooling derivative acquires a* ratification channel *term:*

$$\frac{dP}{d\pi} = u'(\pi)\left[f(u) - f(\theta^*)\,\frac{\partial \theta^*}{\partial u} \right] - f(\theta^*)\,\frac{\partial \theta^*}{\partial r}\, r'(\pi),$$

where the first bracketed term is the standard bar–boundary effect and the second term captures the additional tightening of Stage-2 enforcement.

(ii) *Under the* Litigation-Adjusted Density-Weighted Boundary Responsiveness *condition* (LA-DWBR):

$$\frac{\partial\theta^*}{\partial u} + \frac{\partial\theta^*}{\partial r}\,\frac{r'(\pi)}{u'(\pi)} \geq \frac{f(u)}{f(\theta^*)}, \tag{LA-DWBR}$$

the pooling mass $P(\pi)$ is (weakly) decreasing in π along any pooling corridor.

(iii) Uniform prior + identity beliefs corollary. *If f is uniform and $\mu(m) = m$ (so $f(u) = f(\theta^*)$ and $m^{\text{req}}(u) = u$), condition (LA-DWBR) simplifies to:*

$$\frac{L_r(g^*, r, \sigma)\,r'(\pi)}{u'(\pi)} \geq \alpha + B'(\theta^*).$$

That is, the marginal tightening of Stage-2 enforcement from ratification must dominate the marginal prize slope $\alpha + B'(\theta^*)$.

Proof. We prove each part in turn.

Part (i): decomposition of $dP/d\pi$. Recall the lying-profit function $\Lambda(\theta; u, r) := \alpha\theta + B(\theta) - C\big(m^{\text{req}}(u), \theta\big)$, where $C(m, \theta) \equiv L(g, r, \sigma)$ with $g = m - \theta$. The boundary θ^* is defined implicitly by $\Lambda(\theta^*; u, r) = 0$. We compute its response to u and r separately.

Response to u. By Lemma 5, Λ is C^1 in a neighborhood of (θ^*, u) with $\Lambda_\theta > 0$, and the implicit function theorem gives

$$\frac{\partial\theta^*}{\partial u} = -\frac{\Lambda_u}{\Lambda_\theta}.$$

The numerator is $\Lambda_u(\theta; u) = -C_m\big(m^{\text{req}}(u), \theta\big)\,(\mu^{-1})'(u)$, and on the up-wedge ($m^{\text{req}}(u) > \theta^*$) we have $C_m = \eta + \kappa\Delta > 0$ by (S2)–(S4). The denominator is $\Lambda_\theta = \alpha + B'(\theta^*) + C_m > 0$ as shown in Lemma 4. Hence

$$\frac{\partial\theta^*}{\partial u} = \frac{C_m\big(m^{\text{req}}(u), \theta^*\big)\,(\mu^{-1})'(u)}{\alpha + B'(\theta^*) + C_m\big(m^{\text{req}}(u), \theta^*\big)} \in (0, 1), \tag{2}$$

where the upper bound $\partial\theta^*/\partial u < 1$ follows because the denominator exceeds the numerator (the extra $\alpha + B' > 0$ in the denominator).

The bound $\partial\theta^*/\partial u < 1$ is the fundamental reason why the one-stage model predicts that pooling increases with the threshold. When u rises by one unit, the bar moves up by one unit, but the boundary θ^* moves up by less than one unit—the bar outruns the boundary, widening the pooling interval. The economic reason is a *dampening* in the boundary's response. When the bar rises, the required message m_{req} rises with it, which increases the gap at the boundary and hence the enforcement cost of pooling. This pushes θ^* upward (the C_m in the numerator). But at the boundary, the controller is exactly indifferent between pooling and separating, and the prize from clearing the bar—$\alpha\theta^* + B(\theta^*)$—is pinned to the boundary type's own quality, not the bar's height. The prize slope $\alpha + B'(\theta^*)$ absorbs part of the bar increase without translating it into boundary movement, because a higher bar does not make approval more valuable to the marginal type. This dampening (the $\alpha + B'$ wedge in the denominator) ensures the boundary always lags the bar in the one-stage model. Breaking this result requires a second channel that pushes θ^* upward independently of u—which is exactly what the ratification channel provides.

Response to r. The enforcement cost $C = L(g, r, \sigma)$ now depends on r through Assumption 2. Differentiating $\Lambda(\theta; u, r)$ with respect to r:

$$\Lambda_r(\theta; u, r) \;=\; -L_r(g, r, \sigma) \;\leq\; 0$$

by (L2) (since $g = m^{\text{req}}(u) - \theta^* > 0$ on the up-wedge). Applying the implicit function theorem to $\Lambda(\theta^*; u, r) = 0$:

$$\frac{\partial \theta^*}{\partial r} \;=\; -\frac{\Lambda_r}{\Lambda_\theta} \;=\; \frac{L_r(g^*, r, \sigma)}{\alpha + B'(\theta^*) + C_m\big(m^{\text{req}}(u), \theta^*\big)} \;\geq\; 0. \tag{3}$$

The inequality is strict whenever $L_r(g^*, r, \sigma) > 0$, which holds for $g^* > 0$ by (L2). The economic content is that harder ratification (higher r) raises the marginal cost of lying at every gap, pushing the boundary θ^* upward—more types separate.

Differentiation of $P(\pi)$. The pooling mass is

$$P(\pi) \;=\; \int_{\theta^*(u(\pi),\, r(\pi))}^{u(\pi)} f(\theta)\, d\theta.$$

By Leibniz's rule and the chain rule, with $\theta^* = \theta^*(u(\pi), r(\pi))$:

$$\begin{aligned}
\frac{dP}{d\pi} &= \frac{\partial}{\partial u}\left[\int_{\theta^*}^{u} f(\theta)\, d\theta\right] u'(\pi) \;+\; \frac{\partial}{\partial r}\left[\int_{\theta^*}^{u} f(\theta)\, d\theta\right] r'(\pi) \\
&= \left[f(u) - f(\theta^*)\,\frac{\partial \theta^*}{\partial u}\right] u'(\pi) \;-\; f(\theta^*)\,\frac{\partial \theta^*}{\partial r}\, r'(\pi) \\
&= u'(\pi)\left[f(u) - f(\theta^*)\,\frac{\partial \theta^*}{\partial u}\right] \;-\; f(\theta^*)\,\frac{\partial \theta^*}{\partial r}\, r'(\pi),
\end{aligned}$$

where the first term is the standard bar–boundary effect and the second term is the ratification-channel effect. This establishes part (i).

Part (ii): LA-DWBR condition. Suppose $dP/d\pi \leq 0$. Since $f(\theta^*) > 0$ (interior support) and $u'(\pi) > 0$ (Proposition 15(i)), we may divide both sides of $dP/d\pi \leq 0$ by $f(\theta^*)\, u'(\pi) > 0$ without changing the inequality direction:

$$\frac{f(u)}{f(\theta^*)} \;-\; \frac{\partial \theta^*}{\partial u} \;-\; \frac{\partial \theta^*}{\partial r}\,\frac{r'(\pi)}{u'(\pi)} \;\leq\; 0.$$

Rearranging:

$$\frac{\partial \theta^*}{\partial u} \;+\; \frac{\partial \theta^*}{\partial r}\,\frac{r'(\pi)}{u'(\pi)} \;\geq\; \frac{f(u)}{f(\theta^*)},$$

which is exactly (LA-DWBR). Conversely, if (LA-DWBR) holds, the same algebra in reverse gives $dP/d\pi \leq 0$.

Part (iii): uniform prior and identity beliefs. Assume f is uniform on $[a, b]$ (so $f(\theta) = 1/(b-a)$ for all $\theta \in [a, b]$) and $\mu(m) = m$ (identity beliefs, so $m^{\text{req}}(u) = u$ and $(\mu^{-1})'(u) = 1$). Then:

(a) $f(u)/f(\theta^*) = 1$;
(b) From (2), $\partial\theta^*/\partial u = C_m/(\alpha + B' + C_m) \in (0, 1)$, since the denominator exceeds the numerator by $\alpha + B' > 0$.

The standard (no-litigation) DWBR condition requires $\partial\theta^*/\partial u \geq f(u)/f(\theta^*) = 1$, which fails because $\partial\theta^*/\partial u < 1$. However, the LA-DWBR condition (LA-DWBR) requires only

$$\frac{C_m}{\alpha + B' + C_m} + \frac{L_r(g^*, r, \sigma)}{\alpha + B' + C_m} \cdot \frac{r'(\pi)}{u'(\pi)} \geq 1,$$

where we used (2) and (3) with $(\mu^{-1})'(u) = 1$. Subtracting the first term from both sides:

$$\frac{L_r(g^*, r, \sigma)\, r'(\pi)}{(\alpha + B' + C_m)\, u'(\pi)} \geq 1 - \frac{C_m}{\alpha + B' + C_m} = \frac{\alpha + B'}{\alpha + B' + C_m}.$$

Multiplying both sides by $(\alpha + B' + C_m) > 0$:

$$\frac{L_r(g^*, r, \sigma)\, r'(\pi)}{u'(\pi)} \geq \alpha + B'(\theta^*).$$

This is the stated condition: the marginal tightening of Stage-2 enforcement from ratification must dominate the marginal prize slope. □

The central tension in the result is that the approval threshold does two jobs simultaneously in the coupled regime. First, it sets the bar for persuasion: how good must the controller's disclosure look to win the vote? Second, it sets the bar for peace: how expensive is it to extinguish litigation after the vote? In the one-stage model, the threshold does only the first job. The benchmark prediction—pooling increases with the threshold—follows because a higher bar sweeps more marginal types into the pool, and the dampening force $\alpha + B'(\theta^*)$ in the denominator of $\partial\theta^*/\partial u$ ensures the boundary always lags the bar ($\partial\theta^*/\partial u < 1$). The two-stage model adds the second job, and the two forces push in opposite directions. The persuasion bar expands the set of types whose true quality falls short, creating more potential exaggerators; the peace bar raises the cost of exaggerating, deterring marginal types from doing so. The ratification channel overcomes the dampening force by providing a second, independent lever on θ^*: even though a one-unit increase in u moves the boundary by less than one unit (because the prize slope absorbs part of the increase), the simultaneous increase in r pushes the boundary further upward through $\partial\theta^*/\partial r > 0$, which does not suffer from the same dampening because tougher ratification raises the *cost* of the marginal lie without raising its *prize*. The combined movement $\theta^*_u\, u' + \theta^*_r\, r'$ can exceed u', so the boundary outruns the bar and the pooling interval shrinks. This is the content of LA-DWBR: it holds when the ratification-channel push ($\theta^*_r\, r'/u'$) is large enough to compensate for the shortfall ($f(u)/f(\theta^*) - \theta^*_u$) that the dampening force creates. Convexity of enforcement costs ($\kappa > 0$) amplifies the ratification channel, because bottom-of-pool types—those with the largest gaps—bear disproportionately high costs, so the types removed by coupling are precisely those whose deterrence yields the largest per-type enforcement saving. This reversal has no counterpart in any one-stage signaling model: it requires neither sluggish beliefs nor thin tails, only that the litigation/ratification technology be sufficiently sensitive to the ratification threshold.

Remark. When $r'(\pi) = 0$ (ratification threshold decoupled from the approval rule), the second term in the $dP/d\pi$ formula vanishes and LA-DWBR reduces to the standard GDWBR condition (Assumption 10 in Appendix C). The carry-forward baseline $r(\pi) = \pi$ with $r'(\pi) = 1$ represents the polar case in which the litigation channel is maximally active. As shown in Section **??**, allowing $\sigma = \sigma(g)$ with $\sigma'(g) \geq 0$ adds the term $L_\sigma\, \sigma'(g)$ to the marginal enforcement cost of lying, which strengthens the litigation channel and relaxes the parameter conditions needed for LA-DWBR.

Corollary 1 (Sharp LA-DWBR under ratification microfoundation)**.** *Maintain the hypotheses of Proposition 3 and adopt the ratification-vote microfoundation of Lemma 1 with general resistance distribution G and recovery $R(g, x) = \ell(g) \cdot x$, on the feasible region $r(\pi) \leq 1 - \sigma$. Under the carry-forward baseline $r(\pi) = \pi$ (so $r'(\pi) = 1$), uniform prior, and identity beliefs, the LA-DWBR condition* (LA-DWBR) *becomes:*

$$\frac{p_r(g^*, r)\, C(R^*, r, \sigma) \;+\; p(g^*, r)\, \dfrac{R^* + v^*(r, \sigma)}{1 - \sigma}}{u'(\pi)} \;\geq\; \alpha + B'(\theta^*), \tag{4}$$

where $g^* = m^{\mathrm{req}}(u) - \theta^*$, $R^* := R(g^*, x^*)$, $v^*(r, \sigma) := G^{-1}\big(r/(1-\sigma)\big)$ is the vote-buying cutoff, and C is the ratification cost from Lemma 1. If $r > 1 - \sigma$, cleansing is infeasible in this microfoundation and the effective enforcement burden is unbounded.

Under uniform resistance $G(v) = (1 + v)/2$ on $[-1, 1]$ (Corollary 5), $v^* = 2r/(1 - \sigma) - 1$ and $R^* + v^* = \big(2r - (1-\sigma)(1-R^*)\big)/(1-\sigma)$, so the second term in the numerator reduces to $\big(2r - (1 - \sigma)(1 - R^*)\big)_+/(1 - \sigma)^2$.

Proof. Proof in Appendix D. □

The left side of (4) decomposes the enforcement tightening into two channels. The first, $p_r\, C$, is the *extensive margin*: a tougher threshold raises the probability of suit. The second, $p\, C_r = p\, (R^* + v^*)/(1 - \sigma)$, is the *intensive margin*: conditional on suit, assembling the supermajority costs more because the pivotal voter's reservation price is higher—a distribution-free term from Lemma 1. The right side is the dampening force $\alpha + B'(\theta^*)$. Reversal requires both channels jointly to overcome this dampening; larger r and σ strengthen both.

Figure 2: Pooling mass under decoupled and coupled regimes

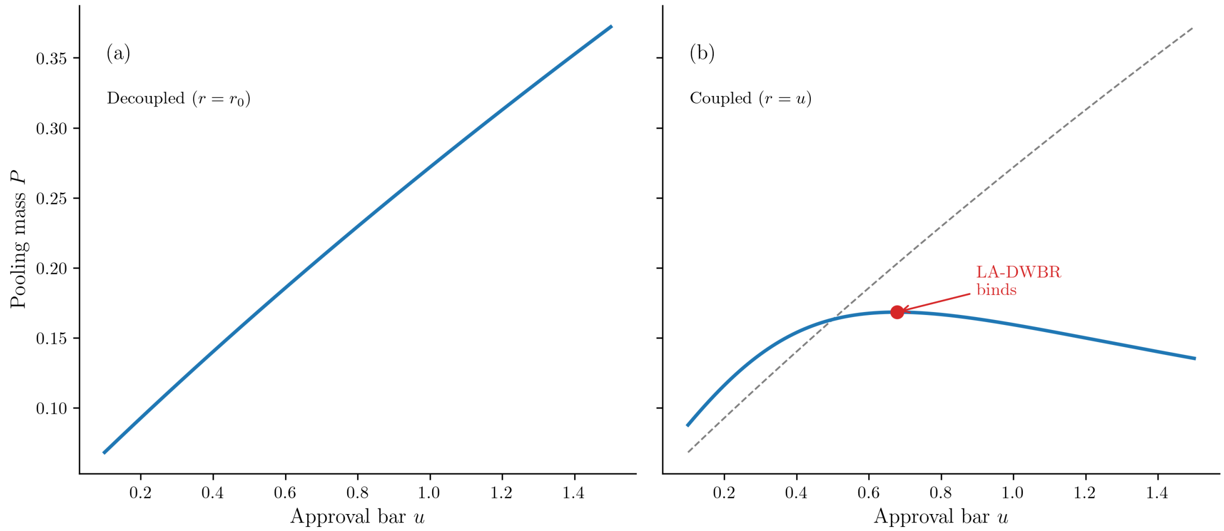

Panel (a): under the decoupled regime ($r = r_0 = 0.5$ fixed), pooling mass $P(u)$ is monotonically increasing in u—the benchmark result of Proposition 16. Panel (b): under the coupled regime ($r = u$), $P(u)$ initially rises but then falls once the LA-DWBR condition (Proposition 3) binds (red dot). The dashed grey line reproduces the decoupled benchmark for comparison. The reversal occurs because the enforcement channel—harder ratification at higher u—deters enough marginal exaggerators to reduce total pooling despite the higher bar. Uniform prior on $[-1, 2]$, identity beliefs, quadratic enforcement costs.

6 Welfare, institutional design, and comparative implications

The preceding analysis shows that coupling the ratification threshold to the approval threshold can reduce the mass of exaggerating controllers. But does reducing pooling improve welfare? Not necessarily. Tighter thresholds deter some value-destroying deals (false positives) but may also deter value-creating deals that now fall below the higher bar (false negatives). Coupling also raises enforcement costs for surviving poolers (intensity effect). Whether the mass reduction dominates these countervailing forces depends on four institutional parameters that the welfare analysis identifies.

The welfare analysis identifies four institutionally interpretable parameters that jointly determine optimal governance design: the deadweight share of enforcement costs (ω), the convexity of enforcement costs in the misrepresentation gap (κ), the sensitivity of enforcement costs to the ratification threshold (L_r), and the quality of the marginal deterred deal (θ^*). The first three govern the enforcement margin—how coupling affects the total cost of policing exaggeration. The fourth governs the selection margin—whether deterring marginal controllers improves or harms firm value. This section derives a welfare decomposition that makes these four parameters the sufficient statistics for the coupling question, and then maps them onto the institutional landscape of ten jurisdictions.

This section develops the welfare implications of the equilibrium objects characterized in Sections 3–5. The core game is unchanged. For each policy parameter π, take as given the induced continuation objects:

$$\big(u(\pi),\ r(\pi),\ \theta^*(u(\pi), r(\pi)),\ m^{\text{req}}(u(\pi))\big).$$

6.1 Welfare criterion

Assumption 4 (W1: Type distribution and integrability)**.** *θ has density f on $[\underline{\theta}_0, \overline{\theta}_0]$ with $\underline{\theta}_0 < 0 < \overline{\theta}_0$, $f(\theta) \geq 0$, $\int_{\underline{\theta}_0}^{\overline{\theta}_0} f(\theta)\, d\theta = 1$, and $\int_{\underline{\theta}_0}^{\overline{\theta}_0} |\theta| f(\theta)\, d\theta < \infty$.*

Assumption 5 (W2: Enforcement regularity and feasible policy domain)**.** *$L(g, r, \sigma)$ is continuous on $\{g \geq 0\} \times \mathcal{R} \times \{\sigma\}$ and continuously differentiable in (g, r) on the interior. Welfare analysis is conducted on a policy region Π^{int} such that $r(\pi) \in \mathcal{R}$ for every $\pi \in \Pi^{\text{int}}$. Under the ratification-vote microfoundation, $\mathcal{R} = (0, 1 - \sigma]$.*

Assumption 6 (W3: Regularity of equilibrium mappings on Π^{int})**.** *On Π^{int}, the policy maps $\pi \mapsto u(\pi)$ and $\pi \mapsto r(\pi)$ are C^1, $\pi \mapsto m^{\text{req}}(u(\pi))$ is C^1, and $\pi \mapsto \theta^*(\pi)$ with*

$$\theta^*(\pi) := \theta^*(u(\pi), r(\pi))$$

is C^1 and satisfies interior partial pooling $\underline{\theta}_0 < \theta^*(\pi) < u(\pi) < \overline{\theta}_0$.

Define the approval and pooling sets

$$\mathcal{A}(\pi) := [\theta^*(\pi), \overline{\theta}_0], \qquad \mathcal{P}(\pi) := [\theta^*(\pi), u(\pi)).$$

Define the equilibrium misrepresentation gap

$$g(\theta; \pi) := \big(m^{\text{req}}(u(\pi)) - \theta\big)\mathbf{1}\{\theta \in \mathcal{P}(\pi)\},$$

so $g(\theta; \pi) \geq 0$ on $\mathcal{P}(\pi)$ and $g(\theta; \pi) = 0$ outside $\mathcal{P}(\pi)$.

Define expected enforcement burden:

$$\text{EC}(\pi) := \int_{\theta^*(\pi)}^{u(\pi)} L(m^{\text{req}}(u(\pi)) - \theta,\ r(\pi),\ \sigma)\, f(\theta)\, d\theta.$$

The two welfare criteria are:

$$W^{\text{firm}}(\pi) := \text{E}[\mathbf{1}\{\theta \in \mathcal{A}(\pi)\}\theta] - \omega\, \text{EC}(\pi),$$

$$W^{\text{tot}}(\pi) := \text{E}\big[\mathbf{1}\{\theta \in \mathcal{A}(\pi)\}\big(\theta + b_0\big)\big] - \omega\, \text{EC}(\pi).$$

Lemma 2 (Exact relation between welfare criteria)**.** *Under Assumptions 4-6,*

$$W^{\text{tot}}(\pi) - W^{\text{firm}}(\pi) = b_0 \Pr\big(\theta \geq \theta^*(\pi)\big).$$

Proof. Subtract the two definitions. The enforcement terms cancel exactly, leaving

$$W^{\text{tot}}(\pi) - W^{\text{firm}}(\pi) = \mathrm{E}[\mathbf{1}\{\theta \geq \theta^*(\pi)\}b_0] = b_0\ \Pr(\theta \geq \theta^*(\pi)).$$

□

Remark. Lemma 2 implies the two criteria are not generically identical in comparative statics: they coincide only when $b_0 = 0$ or when comparing policies with the same approval probability.

6.2 Welfare decomposition: selection errors vs. enforcement frictions

Take the firm-value benchmark first-best cutoff

$$\theta_{\text{FB}} := 0, \qquad W^{\text{FB}} := \int_0^{\overline{\theta}_0} \theta f(\theta)\, d\theta.$$

Define selection-loss terms

$$\text{FP}(\pi) := \mathbf{1}\{\theta^*(\pi) < 0\} \int_{\theta^*(\pi)}^{0} (-\theta) f(\theta)\, d\theta,$$

$$\text{FN}(\pi) := \mathbf{1}\{\theta^*(\pi) > 0\} \int_{0}^{\theta^*(\pi)} \theta f(\theta)\, d\theta.$$

Proposition 4 (Welfare decomposition identity)**.** *Under Assumptions 4-6,*

$$W^{\text{firm}}(\pi) = W^{\text{FB}} - \text{FP}(\pi) - \text{FN}(\pi) - \omega\, \text{EC}(\pi).$$

Proof. Proof in Appendix E, §E.1. □

The decomposition says that any governance regime's welfare shortfall relative to first-best comes from exactly three sources, and these three sources exhaust the gap. False positives are the direct value destruction from approving negative-NPV deals—transactions that a fully informed, conflict-free planner would reject. False negatives are the forgone value from rejecting positive-NPV deals—transactions the planner would approve but that fail to clear the bar because the controller's exaggeration is too costly to sustain. Enforcement costs are the pure deadweight loss from the litigation and ratification process—resources burned on vote-buying, legal fees, and judicial proceedings rather than on productive activity. In the first-best world of perfect information and no conflicts, all three terms are zero; the design question is how to minimize their sum. The division among the three terms is governed by the location of θ^* relative to zero, which is itself determined by the primitives of the enforcement technology—the cost parameters $\eta(r, \sigma)$ and $\kappa(r, \sigma)$ that shape how expensive it is to sustain a given misrepresentation gap, and the ratification threshold r that governs how hard it is to buy peace. When coupling raises r alongside π, it pushes θ^* upward through the enforcement channel of Proposition 3: fewer types find it profitable to lie, and the welfare identity shifts weight from the FP and EC terms toward the FN term.

6.3 Why linking r to π changes welfare predictions

Define

$$\phi(\theta,\pi) := L(m^{\text{req}}(u(\pi)) - \theta,\ r(\pi),\ \sigma)\, f(\theta), \quad \theta^* := \theta^*(\pi),\ u := u(\pi).$$

Proposition 5 (Derivative decomposition for enforcement and welfare)**.** *Under Assumptions 4-6, for $\pi \in \Pi^{\text{int}}$:*

$$\frac{d}{d\pi}\text{EC}(\pi) = \phi(u,\pi)u'(\pi) - \phi(\theta^*,\pi)\theta^{*\prime}(\pi) + \int_{\theta^*}^{u} \big[L_g(m^{\text{req}}(u) - \theta, r(\pi), \sigma)\,(m^{\text{req}})'(u)u'(\pi) + L_r(m^{\text{req}}(u) - \theta, r(\pi), \sigma$$

and

$$\frac{d}{d\pi}W^{\text{firm}}(\pi) = -\theta^*(\pi) f(\theta^*(\pi))\,\theta^{*\prime}(\pi) - \omega\,\frac{d}{d\pi}\text{EC}(\pi).$$

If $\theta^*(u,r)$ is C^1 in (u,r), then $\theta^{*\prime}(\pi) = \theta^*_u\, u'(\pi) + \theta^*_r\, r'(\pi)$ and therefore

$$\frac{d}{d\pi}\text{EC}(\pi) = u'(\pi)\,\mathcal{D}_u(\pi) + r'(\pi)\,\mathcal{D}_r(\pi),$$

for explicit $\mathcal{D}_u, \mathcal{D}_r$ given in Appendix E.

Proof. Proof in Appendix E, §E.2. □

The four terms in $d\text{EC}/d\pi$ have distinct economic content. The *top-boundary term* $\phi(u,\pi)\,u'(\pi)$ captures the enforcement cost brought in by the new type that enters the pooling set at the top when the bar rises. Because this type sits at $\theta = u$, its gap is $m^{\text{req}}(u) - u \approx 0$: it is barely exaggerating. By $L(0,r,\sigma) = 0$, this type's enforcement cost is negligible, so the top-boundary contribution washes out. The *bottom-boundary term* $-\phi(\theta^*,\pi)\,\theta^{*\prime}(\pi)$ captures the enforcement cost *removed* when the marginal liar at θ^* is deterred. This type has the *largest* gap in the pooling set—$g^* = m^{\text{req}}(u) - \theta^*$—and therefore, by convexity of L in g, bears the highest per-capita enforcement cost of any pooler. Whenever coupling raises θ^* (i.e. $\theta^{*\prime} > 0$), it is precisely this most expensive type that drops out. The *interior-gap term* $\int L_g \cdot (m^{\text{req}})' u' f\, d\theta$ reflects the fact that a higher bar forces all surviving poolers to send a larger message, widening each type's gap and raising per-type enforcement costs. This force is present even without coupling and is always welfare-harmful. The *coupling term* $\int L_r \cdot r' f\, d\theta$ is the pure effect of tightening the ratification threshold: for every surviving pooler, each unit of gap is now more costly to sustain because the enforcement technology itself has stiffened. This term is identically zero when $r'(\pi) = 0$ and is always welfare-harmful when $r' > 0$ and $L_r > 0$.

The welfare question on the enforcement margin therefore reduces to whether the bottom-boundary saving—removing the most expensive types—dominates the two cost-escalation forces on survivors. The bottom-boundary and interior forces map onto recognizable transactions. The bottom of the pooling set—large-gap, high-cost types—corresponds to deals like the *Dell* freeze-out (2013), where the offered price was roughly twenty percent below judicially determined fair value; exaggeration on that scale generates years of appraisal litigation and expert costs, placing these types on the steep, convex region of L. By contrast, the top of the pool—a routine related-party lease priced marginally below market—involves a tiny gap, negligible litigation risk, and near-zero enforcement cost on the flat part of the cost curve.

Corollary 2 (Coupling channel versus decoupling channel)**.** *Under the assumptions of Proposition 5, if $r'(\pi) = 0$ (decoupled regime) then*

$$\frac{d}{d\pi}\mathrm{EC}(\pi) = u'(\pi)\,\mathcal{D}_u(\pi).$$

If $r'(\pi) \neq 0$ (coupled regime), the additional term $r'(\pi)\mathcal{D}_r(\pi)$ enters exactly.

Proof. Immediate from Proposition 5. □

When the ratification threshold is decoupled from the approval threshold ($r' = 0$), raising the approval standard affects enforcement costs only through the bar channel: a higher required message widens every pooler's gap. When r is coupled to π ($r' > 0$), an additional coupling channel enters through $r'(\pi)\,\mathcal{D}_r(\pi)$. This channel operates through two opposing sub-forces. First, a *mass effect*: because a higher ratification threshold deters lying—the boundary type θ^* rises—fewer types pool, directly reducing the enforcement-cost integral. Second, an *intensity effect*: for the types that still pool, each faces a steeper enforcement cost surface because r is now tougher.

Convexity of L in the gap g is what makes the mass effect powerful, and the intuition is best seen through the skewed cost structure within the pooling interval $[\theta^*, u)$. Since every pooler sends the same message $m^{\mathrm{req}}(u)$, the gap for type θ in the pool is $g(\theta) = m^{\mathrm{req}}(u) - \theta$, which is mechanically decreasing in θ. The boundary type θ^* has the largest gap and therefore the highest per-type enforcement cost; types near u have gaps approaching zero and near-zero enforcement costs. Enforcement waste is concentrated at the bottom of the pool, and convexity amplifies the skew: bottom types are disproportionately expensive, not merely proportionally so. Coupling removes precisely these types first—controllers whose deals are far worse than what they claim, generating years of appraisal litigation and expert costs on the steep, convex region of L—while routine near-market transactions at the top of the pool, sitting on the flat part of the cost curve, are barely affected. The mass effect therefore operates as a targeted filter: it removes the types whose enforcement costs are disproportionately large relative to their share of the pool. When LA-DWBR holds, this targeted removal from the expensive end dominates the per-unit escalation on survivors.

Combining this enforcement-margin analysis with the selection margin yields the full welfare picture. On the *selection margin*, coupling pushes θ^* upward—welfare-improving when $\theta^* < 0$ (removing negative-NPV deals) and welfare-harmful when $\theta^* > 0$ (over-deterring positive-NPV deals). The two margins align when the pool bottom is populated by negative-NPV types with large gaps: coupling simultaneously reduces false positives and enforcement waste, at the cost of only modest false-negative increase. Restrepo's (2021) finding that MoM adoption surged after *MFW* without measurable effects on deal premiums or completion rates is consistent with this alignment—the deterred deals were marginal, near-zero-NPV transactions that were cheap to lose on the selection margin and expensive to sustain on the enforcement margin. The tension arises only when θ^* has crossed zero and coupling is over-deterring genuinely valuable deals while also escalating costs for remaining poolers.

The deadweight parameter ω governs how much enforcement waste matters for welfare. In the pre-*Trulia* merger-litigation environment, where over ninety percent of deals were challenged and the dominant resolution was a disclosure-only settlement that changed nothing about deal terms, enforcement expenditures were largely waste rather than productive monitoring—a high-ω world in which the enforcement-cost savings from coupling carry substantial welfare weight.

Proposition 3 is a sufficient condition under which the enforcement-feedback component in Proposition 5 dominates the bar-only benchmark component along the relevant policy corridor.

6.4 Example: closed-form welfare under uniform priors and identity beliefs

The uniform-prior, identity-beliefs specialization is not merely a tractability device—it isolates the mechanism in its starkest form. Under these assumptions, the density ratio $f(u)/f(\theta^*)$ equals one, so the standard DWBR condition always fails: the one-stage model unambiguously predicts that pooling increases with the threshold. Any reversal must therefore come entirely from the Stage 2 enforcement channel, making this environment the hardest test case for the coupling mechanism. The closed-form expressions below make the key welfare tradeoffs algebraically transparent. The pooling width $w = u - \theta^*$ is determined by a quadratic equation whose coefficients are exactly the primitives of the enforcement technology: $\eta(r, \sigma)$ captures the fixed marginal cost of any non-zero misrepresentation, $\kappa(r, \sigma)$ captures the convexity that makes large lies disproportionately expensive, and $\alpha + b_1$ is the prize slope that dampens boundary movement. The enforcement cost integral EC decomposes into a linear term (proportional to w^2, driven by η) and a convex term (proportional to w^3, driven by κ), confirming that the cubic term—which grows fastest as the pool widens—is the dominant welfare-relevant force for wide pools. The closed-form welfare expression $W^{\text{firm}} = (\overline{\theta}_0^2 - \theta^{*2})/(2\Delta_\theta) - \omega\,\text{EC}$ shows that coupling improves welfare precisely when the resulting increase in θ^* (which raises θ^{*2} and thus selection quality) and decrease in w (which shrinks enforcement costs) jointly outweigh any false-negative loss from over-deterrence—the same four-parameter condition (ω, κ, L_r, θ^*) now visible in explicit algebra.

Proposition 6 (Closed-form welfare objects under uniform prior and identity beliefs)**.** *Maintain Proposition 16's environment with*

$$\theta \sim U[\underline{\theta}_0, \overline{\theta}_0], \qquad \mu(m) = m, \qquad B(\theta) = b_0 + b_1\theta, \quad b_1 \in (0, 1 - \alpha),$$

and

$$L(g,r,\sigma) = \eta(r,\sigma)\, g + \frac{\kappa(r,\sigma)}{2} g^2, \qquad g \geq 0.$$

Let u be any interior bar and define $w := u - \theta^*(u,r)$. Then:

$$\frac{\kappa(r,\sigma)}{2} w^2 + \big(\alpha + b_1 + \eta(r,\sigma)\big) w = (\alpha + b_1)u + b_0.$$

If $\kappa(r,\sigma) > 0$:

$$w(u,r) = \frac{-\big(\alpha + b_1 + \eta(r,\sigma)\big) + \sqrt{\big(\alpha + b_1 + \eta(r,\sigma)\big)^2 + 2\kappa(r,\sigma)\big((\alpha + b_1)u + b_0\big)}}{\kappa(r,\sigma)},$$

and $\theta^*(u,r) = u - w(u,r)$. If $\kappa(r,\sigma) = 0$:

$$w(u,r) = \frac{(\alpha + b_1)u + b_0}{\alpha + b_1 + \eta(r,\sigma)}.$$

Let $\Delta_\theta := \overline{\theta}_0 - \underline{\theta}_0$. Then

$$\mathrm{EC}(u,r) = \eta(r,\sigma)\frac{w(u,r)^2}{2\Delta_\theta} + \kappa(r,\sigma)\frac{w(u,r)^3}{6\Delta_\theta},$$

and

$$W^{\text{firm}}(u,r) = \frac{\overline{\theta}_0^{\,2} - \theta^*(u,r)^2}{2\Delta_\theta} - \omega\, \mathrm{EC}(u,r).$$

Proof. Proof in Appendix E, §E.3. □

Figure 3: Welfare loss decomposition under decoupled and coupled regimes

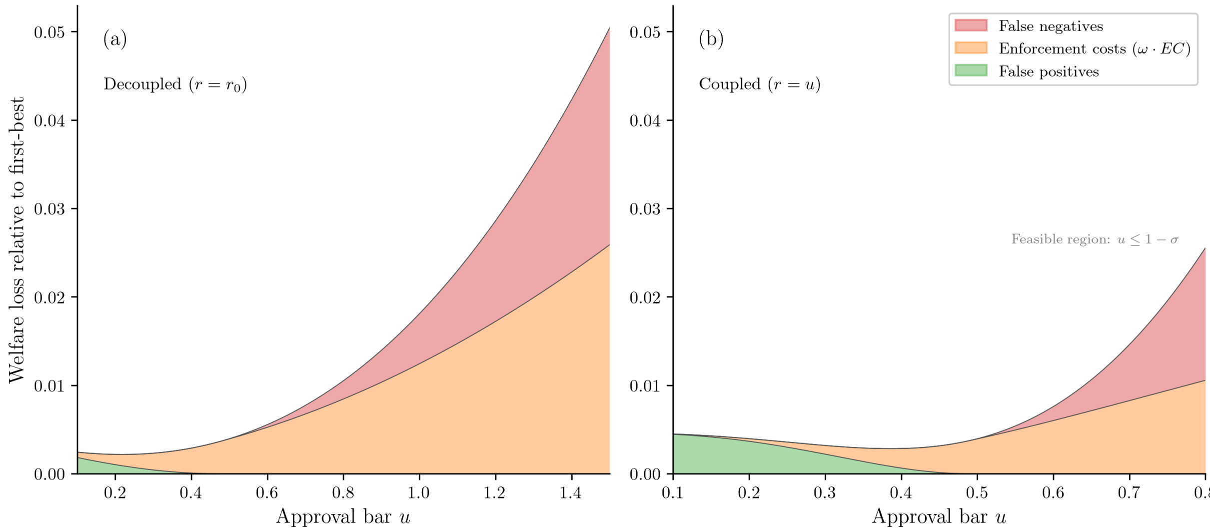

Panel (a): decoupled regime ($r = r_0 = 0.5$). As the bar rises, enforcement costs (orange) grow monotonically because more types pool with larger gaps. Panel (b): coupled regime ($r = u$), plotted over the feasible region $u \leq 1 - \sigma = 0.8$. Enforcement costs are smaller at comparable bar levels—coupling deters the high-gap types that drive enforcement waste (mass effect). False negatives (red) grow more steeply under coupling because the higher enforcement costs deter some value-creating deals (over-deterrence). False positives (green) appear only at low bars where $\theta^* < 0$. The stacked areas sum to the total welfare loss relative to first-best. Shared y-axis scales permit direct cross-panel comparison.

6.5 Comparative institutional implications

The four welfare parameters ($\omega, \kappa, L_r, \theta^*$) are institutionally interpretable features of legal systems (cf. Djankov et al., 2008; Dyck and Zingales, 2004; Ouyang and Zhu, 2016; Shleifer and Vishny, 1997; Atanasov et al., 2011), and different jurisdictions occupy different points in this space. The sign of θ^* is not estimated from a single structural model; it is inferred from each regime's own revealed policy diagnosis: jurisdictions that have tightened controller-transaction requirements (India, Hong Kong) reveal a policy belief that marginal deterred deals are value-destroying ($\theta^* < 0$); those that have loosened requirements (Japan, UK post-2024) reveal a belief that marginal deterred deals are value-creating ($\theta^* > 0$). The comparative exercise maps these stated policy rationales into the model's parameter space—not as conclusions derived from the model itself.

In jurisdictions with concentrated ownership and documented tunneling, the selection and enforcement margins point in the same direction. In India, SEBI mandates majority-of-the-minority approval for material related-party transactions (Li, 2021; Fried et al., 2020; Dai et al., 2019), and promoter groups cannot vote on their own deals, making the procedural hurdle a genuine substantive barrier (L_r high). The case for coupling rests on the selection margin: the marginal deterred deals are value-destroying tunneling transactions ($\theta^* < 0$). Hong Kong presents a similar picture (Kim, 2019; Lim, 2019): Chapter 14A of the HKEX Listing Rules requires an independent financial adviser and a disinterested shareholder vote for connected transactions, and empirical studies document target-firm trading discounts of ten to twenty percent attributable to expropriation risk. In both jurisdictions enforcement costs are low to moderate (ω low–moderate, κ concave), so false-positive reduction rather than enforcement savings drives the welfare case.

China adds a distinctive twist (Chen et al., 2013; Zeng, 2025; Jian and Wong, 2010). The revised PRC Securities Law (2020) and Company Law (2023) introduced statutory penalty multipliers of one to ten times illicit proceeds once the CSRC opens a formal administrative penalty, creating a kinked-convex cost structure (κ step/convex). The mass effect operates with particular force: removing bottom-of-pool types whose large gaps place them above the CSRC trigger generates outsized enforcement savings. Combined with clearly negative θ^* in a market where over thirty percent of listed firms have faced capital occupation by controllers, the model predicts coupling is welfare-improving.

In jurisdictions that have diagnosed over-deterrence, the selection margin reverses. Japan's METI has actively encouraged parent-subsidiary consolidation (Kim, 2019), issuing Fair M&A Guidelines (2019) and Guidelines for Corporate Takeovers (2023) that establish procedural safe harbors. The procedural sensitivity L_r is high, but the marginal deterred deal is likely value-creating ($\theta^* > 0$): coupling would over-deter exactly the transactions the state is trying to promote. The United Kingdom executed a parallel deregulation (Davies, 2019, 2020; Becht et al., 2016): in July 2024 the FCA abolished the mandatory independent shareholder vote for related-party transactions, explicitly motivated by the assessment that the old regime was deterring efficient transactions ($\theta^* > 0$, L_r should be reduced).

Germany's *Konzernrecht* represents an inherently decoupled architecture (Tröger, 2019; Croci et al., 2017; Conac et al., 2007). Under §311 AktG, a controlling enterprise may execute a disadvantageous transaction provided it compensates the subsidiary by year-end; the minority's remedy is an ex post compensation claim via *Spruchverfahren*. Enforcement costs are concave (capped, no punitive damages) and L_r is low by design, implicitly assuming $\theta^* \approx 0$.

The sharpest comparative test is the India–Japan contrast. Both employ majority-of-the-minority voting with special committee oversight and high L_r, but θ^* differs in sign: negative in India (promoter tunneling), positive in Japan (efficiency-enhancing consolidation) (cf. Goshen and Hamdani, 2016). The model predicts that the identical procedural tool is welfare-improving in India and welfare-harmful in Japan—a testable implication for marginal-transaction returns and deal volume.

Delaware is the most contested case: three of four parameters favor coupling. Enforcement costs have a high deadweight share (ω high): disclosure-only settlements dominated merger litigation until *In re Trulia* (2016). The cost structure is step-function with convexity above the entire-fairness trigger (κ step/convex), and the sensitivity of enforcement to the cleansing threshold was very high (L_r very high) under the *MFW* framework.

Yet S.B. 21 (March 2025) moved sharply toward decoupling, allowing non-freeze-out related-party transactions to achieve safe harbor with *either* an independent committee *or* an MoM vote. The welfare justification would require $\theta^* > 0$, but the available evidence does not support that assessment: Restrepo (2021) found that MoM adoption surged after *MFW* without measurable effects on deal premiums or completion rates—consistent with $\theta^* \approx 0$ (see also Gözlügöl, 2022; Boone et al., 2018). This is more consistent with the model's alternative explanation: S.B. 21 reflects controller rent-seeking rather than a welfare-improving correction of over-deterrence (cf. Khoo and Tallarita, 2025). The model provides a diagnostic: if $\theta^* \leq 0$ and (ω, κ, L_r) favor coupling, a reform that decouples is evidence of the governance externality formalized in the private-ordering extension (§6.6.1).

France operates a third regime outside the model's current scope. Under the *conventions réglementées* framework (see Conac et al., 2007; Bianchi et al., 2019; Enriques and Volpin, 2007) (Code de commerce, Art. L225-38 to L225-42), the controller may execute the transaction regardless of minority approval but bears strict liability for prejudicial consequences—an "execute and compensate" architecture that shifts the policy instrument from deterrence to ex post redistribution. Extending the model to accommodate such liability-shifting regimes is a natural direction for future work.

6.6 Extensions: private ordering and committee channels

This section keeps Stages 0-2 unchanged and adds formal Stage -1 and Stage 0.5 extensions. The objective is to map private ordering and committee design into the same equilibrium objects used in the core model.

6.6.1 Private ordering: Stage -1 charter choice

At Stage -1, the controller chooses a governance regime

$$G \in \mathcal{G} := \{L, H\}.$$

Regime L is a lax/decoupled package and regime H is a protective/coupled package. Formally, each G induces continuation primitives

$$\Gamma(G) := (\pi_G, r_G, \sigma_G),$$

with $r_L = r_0$, $r_H = \pi_H$, and $\pi_H \geq \pi_L$.

For each G, let $(u_G, \theta^*_G, m^{\text{req}}_G)$ denote the continuation equilibrium objects from Sections 5.2.3 and 5.3.2, where $m^{\text{req}}_G := m^{\text{req}}(u_G)$. Define

$$I_G(\theta) := \mathbf{1}\{\theta \geq \theta^*_G\}, \qquad g_G(\theta) := \big(m_G(\theta) - \theta\big)_+,$$

with equilibrium message policy

$$m_G(\theta) = \begin{cases} \theta, & \theta < \theta^*_G, \\ m^{\text{req}}_G, & \theta \in [\theta^*_G, u_G), \\ \theta, & \theta \geq u_G. \end{cases}$$

Define pooling mass and expected enforcement burden:

$$P(G) := \int_{\theta_G^*}^{u_G} f(\theta)\, d\theta, \qquad EC(G) := \int_{\theta_G^*}^{u_G} L\big(m_G^{\mathrm{req}} - \theta, r_G, \sigma_G\big) f(\theta)\, d\theta.$$

6.6.2 Private ordering: pricing and controller objective

Let $V_M(G)$ be ex ante minority value under G, normalized per unit minority claim:

$$V_M(G) := \mathrm{E}\big[I_G(\theta)\big(\theta - D(\theta)\big)\big] - \omega\, EC(G),$$

where $\omega \in [0,1]$ is the deadweight share from Section 6.1. Competitive, risk-neutral pricing implies

$$p(G) = V_M(G).$$

Let $\nu \geq 0$ index the level of private-benefit rents and let $\xi \in [0,1]$ be the controller-borne share of enforcement burden. Define

$$\Pi(G;\nu) := \mathrm{E}[I_G(\theta)\, \nu B(\theta)] - \xi\, EC(G),$$

and Stage -1 controller value

$$V_C(G;\alpha,\nu) := \alpha\, p(G) + \Pi(G;\nu).$$

Social value is $W(G)$ from Section 6.

Assumption 7 (Two-regime ordering for Stage -1)**.** *There exist constants $\Delta_M > 0$ and a function $\Delta_\Pi(\nu)$ such that*

$$\Delta_M := V_M(H) - V_M(L) > 0, \qquad \Delta_\Pi(\nu) := \Pi(H;\nu) - \Pi(L;\nu) < 0$$

for relevant ν, and $\Delta_\Pi'(\nu) \leq 0$.

The conditions in Assumption 7 correspond to the $(\omega, \kappa, L_r, \theta^*)$ region where coupling is welfare-improving: $\Delta_M > 0$ requires that false-positive reduction and enforcement savings outweigh false-negative costs, while $\Delta_\Pi < 0$ holds when coupling constrains extraction. Outside this region ($\theta^* > 0$), the assumption fails and the governance externality disappears—controller and planner agree on the lax regime.

The controller trades off a higher share price (capturing $\alpha \cdot \Delta_M$ via competitive pricing) against lost extraction rents ($|\Delta_\Pi|$), adopting protection only when $\alpha \geq \alpha^*(\nu) := -\Delta_\Pi(\nu)/\Delta_M$. The governance externality is therefore *most severe where the welfare case for mandatory coupling is strongest*: high private-benefit intensity simultaneously raises Δ_W and raises $\alpha^*(\nu)$. Markets with concentrated ownership and large extraction rents—India, Hong Kong, China—face the widest gap between socially optimal governance and what private ordering delivers.

Proposition 7 (Cash-flow alignment and regime choice threshold)**.** *Under Assumption 7, define*

$$\alpha^*(v) := -\frac{\Delta_\Pi(v)}{\Delta_M} > 0.$$

Then

$$V_C(H; \alpha, v) - V_C(L; \alpha, v) = \alpha\Delta_M + \Delta_\Pi(v),$$

so the controller chooses the protective regime H if and only if $\alpha \geq \alpha^*(v)$. In particular, the chosen regime is weakly more protective as α increases.

Proof. See Appendix D, §D.1.1. □

Proposition 8 (Private benefits tilt choice toward decoupling)**.** *Under Assumption 7, $\alpha^*(v)$ is weakly increasing in v. Equivalently, for any fixed α, an increase in private-benefit intensity weakly expands the parameter region in which the controller chooses the lax regime L.*

Proof. See Appendix D, §D.1.2. □

Higher v simultaneously raises $\alpha^*(v)$ (more controllers choose the lax regime) and raises Δ_W (coupling is more socially valuable)—a vicious circle. The controllers who would benefit society most from adopting protection are precisely those who will not do it voluntarily, formalizing the case for mandatory procedural standards.

Assumption 8 (Planner ranking)**.**

$$\Delta_W := W(H) - W(L) > 0.$$

Proposition 9 (Governance externality under partial alignment)**.** *Under Assumptions 7 and 8, if $\alpha < \alpha^*(v)$ then the controller chooses L while welfare is maximized by H. Hence*

$$\arg\max_{G\in\{L,H\}} V_C(G; \alpha, v) \neq \arg\max_{G\in\{L,H\}} W(G).$$

Proof. See Appendix D, §D.1.3. □

The controller-borne share of enforcement costs ξ mediates the externality's severity. Under the English Rule (UK, Australia, Singapore), the losing party bears both sides' costs (ξ high), partially correcting the externality. Under the American Rule (ξ lower, with D&O insurance absorbing exposure), the externality is exacerbated. This connects the private ordering model to cross-jurisdictional variation in cost-shifting rules.

Figure 4: Governance externality in (ν, α) space

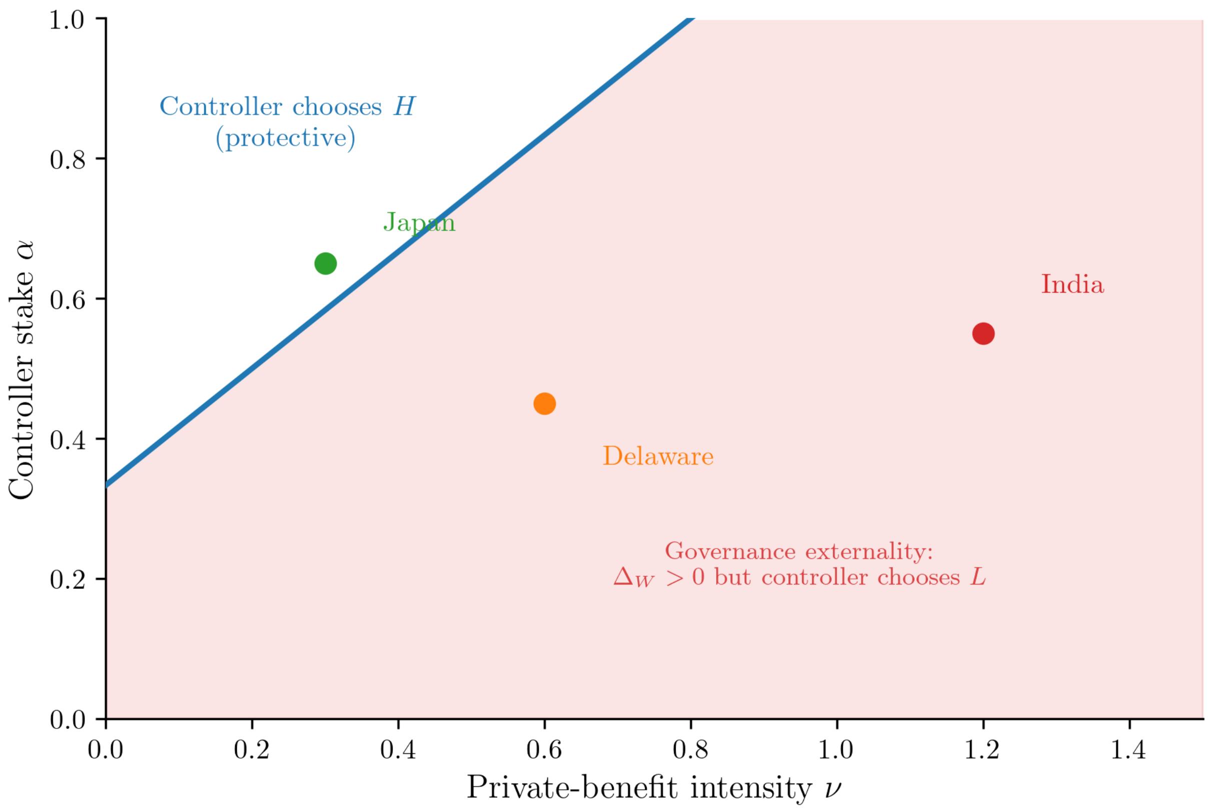

The boundary $\alpha^*(\nu) = -\Delta_\Pi(\nu)/\Delta_M$ (blue line) separates controllers who voluntarily adopt protective governance (above) from those who choose the lax regime (below). The shaded region below the boundary is the governance externality zone: coupling is welfare-improving ($\Delta_W > 0$) but the controller chooses L. The boundary is increasing in ν—higher private-benefit intensity raises the controller's resistance to protection—so the externality is most severe where it matters most (vicious circle). Approximate jurisdiction placements: Japan (green, low ν, high α, above the line—voluntary adoption); Delaware (orange, moderate ν, moderate α, below the line—borderline); India (red, high ν, moderate α, deep in the externality zone). Stylized parameterization for illustration.

SB 21 amended DGCL Section 144 to permit business judgment protection for non-going-private controller transactions (revised §144(b)) with *either* disinterested-stockholder approval *or* independent-committee approval; going-private transactions under §144(c) continue to require both.[8] The either/or structure interacts with Corollary 3: supermodularity implies that SB 21's corner solutions are welfare-inferior to the joint requirement of *MFW*. When $\alpha < \alpha^*(\nu)$, controllers select the procedural prong constraining extraction least (Proposition 9). Under the identified parameter conditions, SB 21 is therefore welfare-inferior on two margins: either/or forfeits procedural complementarity, and private ordering enables controllers to select the welfare-inferior corner. Khoo and Tallarita (2025) document evidence consistent with controller rent-seeking in the reform's legislative history.

[8]Delaware General Assembly, SB21: https://legis.delaware.gov/BillDetail/141930; implementation context: https://www.reuters.com/legal/legalindustry/delaware-law-changes-parameters-transactions-involving-interested-directors-2025-05-08/.

6.6.3 Committee as a second procedural lever

Committees act as a pre-vote filter: a more independent committee shifts the mix of deals reaching the vote toward higher quality, reducing both the mass of exaggerators and expected enforcement costs. The committee is procedurally complementary to vote and ratification design, addressing a different margin (pre-vote composition) while reinforcing the same disciplinary channel (Gilson and Gordon, 2003; Bebchuk and Hamdani, 2017). This extension adds a committee lever to the Stage 0–2 backbone.[9]

Add Stage 0.5 before Stage 1: the committee either allows the proposal to proceed to the shareholder vote or stops it. Let $\iota \in [0,1]$ denote committee independence and let $Q : [\underline{\theta}_0, \overline{\theta}_0] \to [0,1]$ satisfy $Q'(\theta) > 0$. Define the pass probability

$$q_\iota(\theta) := (1-\iota) + \iota Q(\theta) \in (0,1],$$

and the induced type density among deals reaching the vote

$$f_\iota(\theta) := \frac{q_\iota(\theta) f(\theta)}{Z_\iota}, \qquad Z_\iota := \int_{\underline{\theta}_0}^{\overline{\theta}_0} q_\iota(t) f(t)\, dt.$$

Appendix D confirms that higher committee independence shifts the conditional type distribution in the MLR sense, reduces low types reaching the vote, and lowers enforcement costs at any fixed bar. The signal channel is similarly well-behaved. The economically interesting question is not whether independence helps—it does—but whether its benefits interact with enforcement toughness, and at what cost.

The committee observes $s = \theta + \varepsilon$ with $\mathrm{E}[\varepsilon] = 0$ and ε independent of θ. With probability ι, it recommends $y = \mathbf{1}\{s \geq s_0\}$; with probability $1-\iota$, it is captured and sets $y = 1$. The pass probability is $p_\iota(\theta) := 1 - \iota G_\varepsilon(s_0 - \theta)$, and the effective posterior given message m and recommendation y is $\mu_\iota(m,y) = \tau(m,y)\,\bar{\mu}_\iota(m,y) + (1-\tau(m,y))\mu_0$, where $\bar{\mu}_\iota(m,y) := \mathrm{E}[\theta \mid m, y]$. The monotonicity properties of the signal channel—favorable recommendations lower the required persuasive message and weakly reduce enforcement exposure—are established in Appendix D.

6.6.4 Cost of committee independence and procedural complementarity

The preceding results establish that committee independence weakly reduces false positives and enforcement costs at fixed voting and enforcement parameters. But a more independent committee also blocks deals that a captured committee would have let through—including deals with positive NPV. This subsection formalizes that cost, then shows that the benefit of committee independence is amplified by tougher enforcement while the cost is invariant, generating welfare supermodularity in committee independence and ratification toughness.

[9]On *MFW*'s dual-protection framework, see *Kahn v. M&F Worldwide Corp.*, https://law.justia.com/cases/delaware/supreme-court/2014/334-2013.html.

Definition 1 (Gate false-negative cost)**.** *Define the expected value of positive-NPV deals blocked at the gate:*

$$\mathrm{FN}^{\mathrm{gate}}(\iota) := \int_0^{\overline{\theta}_0} \big(1 - q_\iota(\theta)\big)\, \theta\, f(\theta)\, d\theta.$$

Proposition 10 (Gate false-negative cost is increasing in independence)**.** $\mathrm{FN}^{\mathrm{gate}}(\iota)$ *is strictly increasing and linear in* ι *on* $[0,1]$*, with* $\mathrm{FN}^{\mathrm{gate}}(0) = 0$*. Moreover,* $\mathrm{FN}^{\mathrm{gate}}$ *is independent of* (u, r, σ)*.*

Proof. Since $q_\iota(\theta) = (1-\iota) + \iota Q(\theta)$, we have $1 - q_\iota(\theta) = \iota\big(1 - Q(\theta)\big)$. Hence

$$\mathrm{FN}^{\mathrm{gate}}(\iota) = \iota \int_0^{\overline{\theta}_0} \big(1 - Q(\theta)\big)\, \theta\, f(\theta)\, d\theta.$$

The integral is a positive constant (since $Q(\theta) < 1$ for $\theta < \overline{\theta}_0$ and $\theta > 0$ on the domain of integration), so $\mathrm{FN}^{\mathrm{gate}}$ is strictly increasing and linear in ι with $\mathrm{FN}^{\mathrm{gate}}(0) = 0$. No term depends on (u, r, σ). □

A more independent committee kills more bad deals but also blocks some good ones—the cost that the preceding monotonicity results omit. Critically, $\mathrm{FN}^{\mathrm{gate}}$ does not depend on (r, σ): the gate blocks deals before enforcement. The benefit side depends on (r, σ) through L. This asymmetry is the source of the complementarity result below.

Proposition 11 (Enforcement-channel complementarity: cross-partial in (ι, r))**.** *Fix* $(u, \theta^*, m^{\mathrm{req}}, \sigma)$ *with* $\theta^* < u$ *and define*

$$\mathrm{EC}(\iota, r) := \int_{\theta^*}^{u} q_\iota(\theta)\, L\big(m^{\mathrm{req}} - \theta,\ r,\ \sigma\big)\, f(\theta)\, d\theta.$$

If $L_r(g, r, \sigma) \geq 0$ *for all* $g \geq 0$*, then*

$$\frac{\partial^2\, \mathrm{EC}}{\partial \iota\, \partial r} \;\leq\; 0,$$

with strict inequality if $L_r(m^{\mathrm{req}} - \theta, r, \sigma) > 0$ *and* $Q(\theta) < 1$ *on a subset of* $[\theta^*, u)$ *with positive* f*-measure.*

Proof. Differentiating under the integral:

$$\frac{\partial\, \mathrm{EC}}{\partial \iota} = \int_{\theta^*}^{u} \big(Q(\theta) - 1\big)\, L\big(m^{\mathrm{req}} - \theta, r, \sigma\big)\, f(\theta)\, d\theta \;\leq\; 0,$$

since $Q(\theta) \leq 1$ and $L \geq 0$. Differentiating again with respect to r:

$$\frac{\partial^2\, \mathrm{EC}}{\partial \iota\, \partial r} = \int_{\theta^*}^{u} \big(Q(\theta) - 1\big)\, L_r\big(m^{\mathrm{req}} - \theta, r, \sigma\big)\, f(\theta)\, d\theta.$$

Each factor in the integrand satisfies $Q(\theta) - 1 \leq 0$ and $L_r \geq 0$, so the integrand is weakly negative pointwise. Under the stated non-degeneracy condition, the integral is strictly negative. □

The sign of $\partial^2 \mathrm{EC}/\partial\iota\,\partial r \leq 0$ says that the marginal enforcement-cost reduction from a more independent committee is *larger* when the ratification threshold is tougher. Intuitively, when r is high, each remaining pooler bears steep enforcement costs, so removing a type from the pool via the gate saves more per type removed. A tougher enforcement environment amplifies the value of pre-vote screening.

Corollary 3 (Welfare supermodularity in committee independence and ratification toughness)**.** *Under the conditions of Propositions 10 and 11, define the committee-augmented welfare at fixed* $(u, \theta^*, m^{\mathrm{req}}, \sigma)$*:*

$$\mathcal{W}(\iota, r) := \underbrace{\int_{\theta^*}^{\bar{\theta}_0} q_\iota(\theta)\,\theta\, f(\theta)\, d\theta}_{\text{expected approved value}} - \omega\,\mathrm{EC}(\iota, r) - \mathrm{FN}^{\mathrm{gate}}(\iota).$$

Then

$$\frac{\partial^2 \mathcal{W}}{\partial\iota\,\partial r} = -\omega\,\frac{\partial^2\,\mathrm{EC}}{\partial\iota\,\partial r} \geq 0,$$

since $\partial^2 \mathrm{FN}^{\mathrm{gate}}/\partial\iota\,\partial r = 0$ *and* $\partial^2/\partial\iota\,\partial r$ *of the approved-value integral vanishes at fixed* θ^**. Therefore* $\mathcal{W}$ *is supermodular in* (ι, r)*: committee independence and ratification toughness are welfare complements.*

Proof. The gate false-negative cost $\mathrm{FN}^{\mathrm{gate}}(\iota)$ is independent of r (Proposition 10), so its cross-partial vanishes. The approved-value integral at fixed θ^* depends on ι through $q_\iota(\theta)$ but not on r, so its cross-partial also vanishes. The only surviving term is $-\omega\,\partial^2 \mathrm{EC}/\partial\iota\,\partial r \geq 0$ by Proposition 11. □

Remark (Conservativeness). The supermodularity is established at fixed $(u, \theta^*, m^{\mathrm{req}})$. The result is conservative: $\mathrm{FN}^{\mathrm{gate}}$ overstates the true gate cost (types in $(0, \theta^*)$ would be rejected at the vote anyway), and raising r pushes θ^* upward, shrinking the set harmed by gate false negatives. Both adjustments favor supermodularity.

The economic logic of Corollary 3 is an asymmetry: the *cost* of committee independence (blocking good deals) is invariant to enforcement parameters, while the *benefit* (reducing enforcement waste) scales with r through L_r. Tougher enforcement raises the benefit without raising the cost, so the marginal return to committee quality is increasing in enforcement toughness.

Under *MFW*, Delaware required both an independent committee and MoM approval—maximizing ι and r jointly, consistent with supermodularity. S.B. 21's *either/or* safe harbor permits corner solutions, resting on the premise that the two instruments are substitutes. But they address *different* margins—composition (committee) versus discipline (vote/ratification)—and these margins are complements. Supermodularity implies corner solutions are generically welfare-inferior.

Where $\theta^* < 0$ (India, Hong Kong, China), the gate's false-negative cost is small and both forces favor high (ι, r). Where $\theta^* > 0$ (Japan, the UK), the optimal *level* of both instruments is lower, but supermodularity still holds locally: the complementarity governs the *shape* of the welfare surface; the sign of θ^* governs the *location* of the optimum.

Figure 5: Welfare supermodularity over committee independence and ratification toughness

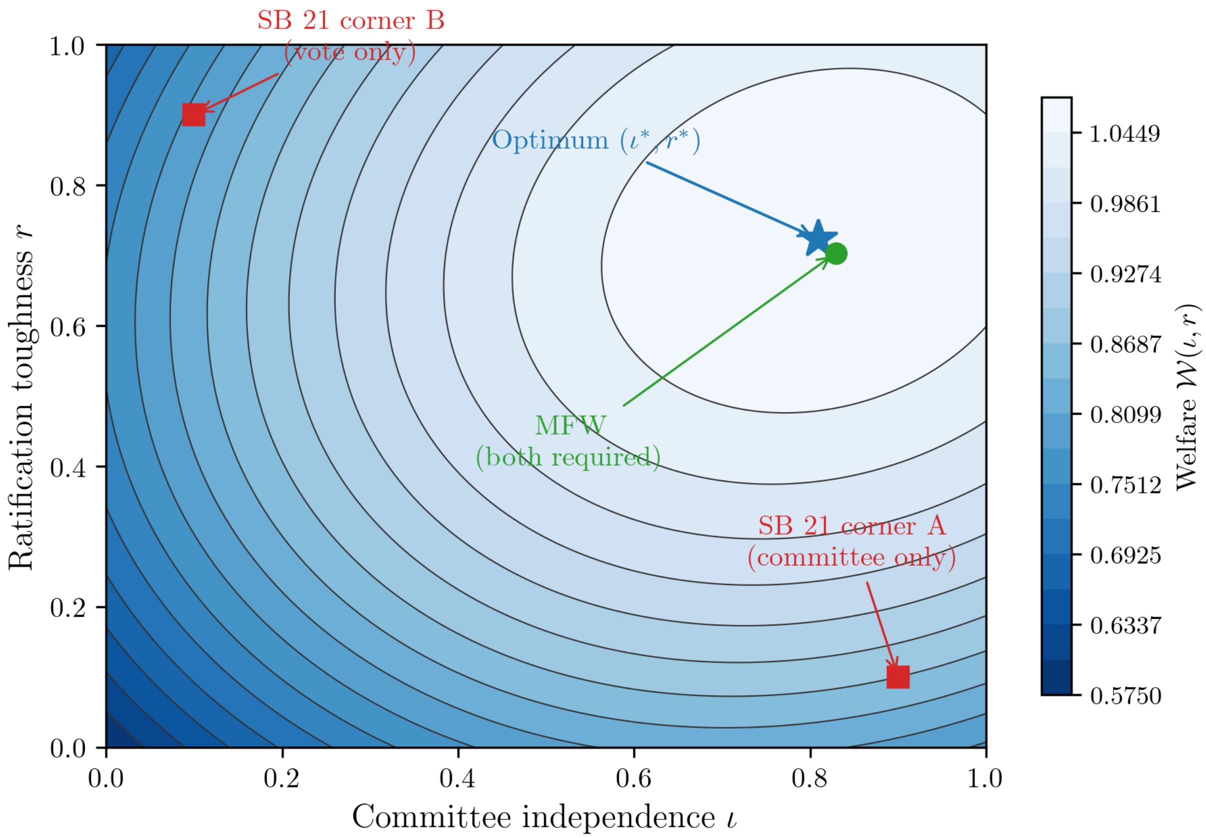

Lighter shading indicates higher welfare $\mathcal{W}(\iota, r)$. The optimum (blue star) is interior, reflecting the complementarity between the two instruments. The MFW point (green circle), which requires both an independent committee and a majority-of-the-minority vote, is near the optimum. The two SB 21 corners (red squares)—committee only (ι high, r low) and vote only (ι low, r high)—are visibly welfare-inferior. The contour lines bend toward the (high, high) corner, reflecting supermodularity: the marginal return to committee independence is increasing in ratification toughness (Corollary 3). Stylized welfare function calibrated to illustrate the qualitative result.

7 Conclusion

This paper started from a puzzle: if tougher shareholder votes should discipline controllers, why does the standard signaling model predict they increase misrepresentation? The answer is that the standard model treats the vote as the end of the story. When the difficulty of buying post-vote peace is linked to the approval threshold, tougher votes raise both the bar for persuasion and the price of peace, and the second channel can dominate.

The welfare implications are conditional. Which force dominates depends on four institutionally interpretable parameters—ω, κ, L_r, and θ^*—and different jurisdictions occupy different regions of this space. Coupling is welfare-improving where marginal deals are value-destroying (India, Hong Kong, China) and welfare-harmful where they are value-creating (Japan, the UK). Delaware is diagnostic: three of four parameters favor coupling, yet S.B. 21 decoupled—consistent with the governance externality rather than efficient design (Khoo and Tallarita, 2025).

The private ordering extension shows that controllers internalize minority-value gains only in proportion to α, producing a governance externality whenever $\alpha < \alpha^*(\nu)$ (Proposition 9). High private-benefit intensity simultaneously raises the welfare gain from coupling and raises the controller's resistance to adopting it—a vicious circle most severe in concentrated-ownership, high-tunneling jurisdictions. The allocation of enforcement costs mediates the externality's severity: the English Rule (ξ high) partially corrects it; the American Rule (ξ low) exacerbates it.

The committee extension establishes that committee independence and ratification toughness are welfare complements (Corollary 3). The complementarity derives from the sequential institutional structure: the *cost* of independence (blocking good deals at the gate) is invariant to the enforcement environment, while the *benefit* (removing high-gap types before they generate enforcement waste) scales with r through L_r. The committee operates on the composition margin; the vote operates on the discipline margin; these margins are complements. Under the model's parameter conditions, S.B. 21's either/or safe harbor is welfare-inferior on two margins: it forfeits procedural complementarity, and the private ordering result predicts controllers will exploit the menu to select the welfare-inferior corner. Extending the model to liability-shifting regimes (French and German architectures) is a natural direction for future work.

A Appendix A: General signaling equilibrium lemmas

A.1 Deferred proof of Proposition 2

Proof. Preliminaries (voting and shape): By (M1), μ is continuous, strictly increasing, and invertible, so $m_{\text{req}}(u)$ exists and is unique and satisfies $\mu(m_{\text{req}}(u)) = u$. Define $\Lambda(\theta; u) = \alpha\theta + B(\theta) - C(m_{\text{req}}(u), \theta)$. Assumptions (S3)-(S4) and (P3) imply:

(a) $\theta \mapsto \Lambda(\theta; u)$ is continuous and strictly increasing because

$$\Lambda_\theta(\theta; u) = \alpha + B'(\theta) - C_\theta\big(m_{\text{req}}(u), \theta\big) \;\geq\; \alpha + B'(\theta) \;>\; 0;$$

(b) by Lemma 3, $\Lambda_u(\theta; u) = -\, C_m(m_{\text{req}}(u), \theta)\, \mu^{-1\prime}(u) \leq 0$, with strict inequality whenever $m_{\text{req}}(u) > \theta$.

(i) Existence. We verify sender optimality type-by-type and then receivers' best responses.

Low types $\theta < \theta^*(u)$. By definition of $\theta^*(u)$, $\Lambda(\theta; u) < 0$. If a low type deviates to the least-cost approved message $m_{\text{req}}(u)$, her payoff is $\Lambda(\theta; u) < 0$, strictly below truthful rejection (0 by (S2)). Any other approved message $m > m_{\text{req}}(u)$ is strictly more costly because $C_m(m, \theta) > 0$ and is increasing on the up-wedge (S3), hence strictly worse. Any under-report $m < \theta$ is rejected since μ is increasing and $m < m_{\text{req}}(u)$ implies $\mu(m) < u$; moreover $C(m, \theta) > C(\theta, \theta) = 0$ by (S2)-(S3). Thus truth maximizes payoff for $\theta < \theta^*(u)$.

Persuaders $\theta \in [\theta^*(u), u)$. Truth would be rejected ($\mu(\theta) < u$), delivering 0. Sending $m_{\text{req}}(u)$ guarantees approval with payoff $\Lambda(\theta; u) \geq 0$. Among approved messages, $m_{\text{req}}(u)$ minimizes signaling cost by the same up-wedge argument as above, hence is strictly optimal.

High types $\theta \geq u$. Under the candidate strategy, high types separate truthfully on $[u, \overline{\theta}_0]$. By Bayes' rule along this separating segment, $\mu(m = \theta) = \theta \geq u$, so truth is approved at zero cost (S2). Any $m > \theta$ raises cost (no incremental benefit), and any $m < \theta$ either remains approved yet adds strictly positive cost on the down-wedge (S3), or falls below u and is rejected. Hence $m = \theta$ is optimal.

Receiver best responses and beliefs. Given the message strategy $m(\cdot)$ above, the on-path messages are (i) $m = \theta$ for $\theta \geq u$ and (ii) $m = m_{\text{req}}(u)$ for $\theta \in [\theta^*(u), u)$. On the separating segment, Bayes' rule gives $\mu(\theta) = \theta \geq u$, so truth is approved. At the pooling message, $m_{\text{req}}(u)$ meets the message threshold by construction (see Remark **??**), so minorities approve.

For off-path messages $m' \notin \{m_{\text{req}}(u)\} \cup [u, \overline{\theta}_0]$, we assign GP-credible beliefs. By the Grossman–Perry criterion, the posterior at any off-path m' is supported on the set of types that would gain from deviating to m' given that belief. For $m' < m_{\text{req}}(u)$ (below the pooling message), any type sending m' is rejected; since truth at $m = \theta < m_{\text{req}}(u)$ is also rejected at zero cost, no type profits from the deviation, so the GP set is empty and beliefs can assign any $\mu(m') < u$. For $m' \in (m_{\text{req}}(u), u)$, the type with the strongest single-crossing incentive to deviate upward is the type $\theta = m'$ (by $C_{m\theta} < 0$ on the up-wedge from (S4)), so GP beliefs assign $\mu(m') = m' < u$ and the message is rejected. No type benefits. Therefore the stated profile is a PSE.

(ii) Allocation uniqueness for fixed u. Let $(\hat{m}, \widehat{\text{beliefs}})$ be any PSE with the same cut-off u. We show the implementation set is $\{\theta \geq \theta^*(u)\}$.

Claim A. No $\theta < \theta^*(u)$ can be implemented. If $\hat{m}(\theta)$ induces approval, then among approved messages the least costly is $m_{\text{req}}(u)$ (by S3), giving payoff $\Lambda(\theta; u) < 0$, strictly below truthful rejection 0.

Claim B. Every $\theta > \theta^*(u)$ must be implemented. We distinguish two sub-cases.

Sub-case B1: $\theta \in (\theta^*(u), u)$. If $\hat{m}(\theta)$ is rejected, deviating to $m_{\text{req}}(u)$ yields payoff $\Lambda(\theta; u) > 0$ (by single-crossing, Lemma 4, and $\theta > \theta^*(u)$), a strict improvement. Among approved messages, $m_{\text{req}}(u)$ minimizes cost on the up-wedge (S3), so $m_{\text{req}}(u)$ is strictly optimal. Hence θ must be implemented.

Sub-case B2: $\theta \geq u$. Truthful reporting $m = \theta$ is approved under GP-credible beliefs: on the separating segment $[u, \overline{\theta}_0]$, Bayes gives $\mu(\theta) = \theta \geq u$. The payoff from truth is $\alpha\theta + B(\theta) - C(\theta, \theta) = \alpha\theta + B(\theta) \geq \alpha u + B(u) \geq 0$, where the first inequality uses $\alpha + B' > 0$ (P3) and the second is (CPC). If instead $\hat{m}(\theta)$ is rejected, payoff is 0, strictly worse. If $\hat{m}(\theta)$ induces approval via some $m \neq \theta$, then $C(m, \theta) > 0 = C(\theta, \theta)$ by (S2), so truth at zero cost strictly dominates.

Thus all $\theta > \theta^*(u)$ are implemented.

At $\theta = \theta^*(u)$, $\Lambda(\theta^*(u); u) = 0$; the type is indifferent between truthful rejection and sending $m_{\text{req}}(u)$, and either tie-break leaves all other allocations unchanged. (In the boundary cases $\theta^*(u) = u$ or $\theta^*(u) = \underline{\theta}_0$, indifference is vacuous: the equilibrium reduces to full separation or full pooling below u, as in part (iii).) Combining the claims gives the stated allocation uniqueness.

(iii) Pooling vs. separation. Set $\Psi_\ell(u) = \Lambda(\underline{\theta}_0; u)$ and $\Psi_h(u) = \lim_{\theta\uparrow u} \Lambda(\theta; u) = \alpha u + B(u) - C(m_{\text{req}}(u), u)$. Since $\Lambda(\cdot; u)$ is continuous and strictly increasing in θ on $[\underline{\theta}_0, u]$, the three cases follow by the intermediate value theorem: $\Psi_\ell(u) < 0 < \Psi_h(u)$ yields a unique $\theta^*(u) \in (\underline{\theta}_0, u)$ (partial pooling); $\Psi_h(u) \le 0$ gives $\theta^*(u) = u$ (full separation); $\Psi_\ell(u) \ge 0$ gives $\theta^*(u) = \underline{\theta}_0$ (full pooling below u). □

A.2 General lemmas

Lemma 3 (Sensitivity of Λ to the rule cut-off)**.** *Fix an interior belief cut-off $u \in (\underline{\theta}, \bar{\theta})$ and let $m^{\text{req}}(u) := \mu^{-1}(u)$. For every (θ, u) with $m^{\text{req}}(u) \neq \theta$,*

$$\frac{\partial \Lambda}{\partial u}(\theta; u) \;=\; -\,C_m\big(m^{\text{req}}(u), \theta\big)\, \mu^{-1\prime}(u).$$

Because $\mu^{-1\prime}(u) > 0$ by (M1)*, the sign is*

$$\operatorname{sign}\Big(\tfrac{\partial \Lambda}{\partial u}\Big) \;=\; -\operatorname{sign}\big(C_m(m^{\text{req}}(u), \theta)\big) \;=\; \operatorname{sign}\big(\theta - m^{\text{req}}(u)\big).$$

In particular, $\frac{\partial \Lambda}{\partial u} < 0$ whenever $m^{\text{req}}(u) > \theta$ and $\frac{\partial \Lambda}{\partial u} > 0$ whenever $m^{\text{req}}(u) < \theta$. At the boundary $m^{\text{req}}(u) = \theta$, the one-sided derivatives in u exist and have the corresponding signs.

Proof. By definition $m^{\text{req}}(u) = \mu^{-1}(u)$, so $\frac{\partial m^{\text{req}}}{\partial u} = \mu^{-1\prime}(u) > 0$ by (M1). Using the definition $\Lambda(\theta; u) = \alpha\theta + B(\theta) - C(m^{\text{req}}(u), \theta)$ and applying the chain rule to the composition $u \mapsto C(m^{\text{req}}(u), \theta)$ (which is C^1 off the diagonal by (S1) when $m^{\text{req}}(u) \neq \theta$) yields

$$\frac{\partial \Lambda}{\partial u}(\theta; u) = \,-\, C_m\big(m^{\text{req}}(u), \theta\big)\, \frac{\partial m^{\text{req}}}{\partial u} \;=\; -\, C_m\big(m^{\text{req}}(u), \theta\big)\, \mu^{-1\prime}(u).$$

Hence $\operatorname{sign}(\partial\Lambda/\partial u) = -\operatorname{sign}\big(C_m(m^{\text{req}}(u), \theta)\big)$. By (S4) and (S2), on the down-wedge ($m < \theta$) we have $C_m(m, \theta) = -\eta + \kappa(m - \theta) < 0$, while on the up-wedge ($m > \theta$) we have $C_m(m, \theta) = +\eta + \kappa(m - \theta) > 0$; the two meet at $m = \theta$ with the one-sided limits $\mp\eta$. Substituting $m = m^{\text{req}}(u)$ gives the stated regional sign: negative when $m^{\text{req}}(u) > \theta$ and positive when $m^{\text{req}}(u) < \theta$.[10] □

Lemma 4 (Single-crossing of the lying-profit function)**.** *Fix any interior belief cut-off $u \in (\underline{\theta}_0, \overline{\theta}_0]$ and define*

$$m^{\text{req}}(u) \;:= \mu^{-1}(u), \qquad \Lambda(\theta; u) \;:= \alpha\,\theta + B(\theta) - C\big(m^{\text{req}}(u), \theta\big).$$

[10] If $m^{\text{req}}(u) = \theta$, C may be kinked along the diagonal (allowed by (S1)); the one-sided derivatives obtained from the two wedges deliver the corresponding signs. Differentiability at the boundary is not needed for the comparative-statics results used later.

Under (P3) *and* (S4) *(with* $\alpha > 0$*), the map* $\theta \mapsto \Lambda(\theta; u)$ *is strictly increasing on* $[\underline{\theta}_0, u]$. *Consequently, for each fixed* u *the equation* $\Lambda(\theta; u) = 0$ *has* at most one *solution* $\theta^*(u) \in [\underline{\theta}_0, u]$. *If in addition there exist* θ_1, θ_2 *with* $\Lambda(\theta_1; u) < 0 < \Lambda(\theta_2; u)$*, then there is a* unique $\theta^*(u) \in (\theta_1, \theta_2)$ *with* $\Lambda(\theta^*(u); u) = 0$.

Proof. By (S1) and (M1), C is C^1 off the diagonal and μ^{-1} is continuous, so Λ is continuously differentiable in θ on $[\underline{\theta}_0, u]$, with

$$\Lambda_\theta(\theta; u) = \alpha + B'(\theta) - C_\theta\big(m^{\text{req}}(u), \theta\big).$$

On the relevant region $\theta \in [\underline{\theta}_0, u]$ we have $m_{\text{req}}(u) \geq \theta$ (up-wedge), hence by (S2)-(S4)

$$C_\theta\big(m_{\text{req}}(u), \theta\big) = -(\eta + \kappa\Delta) \leq -\eta < 0 \quad \text{for } \Delta := m_{\text{req}}(u) - \theta \geq 0.$$

Therefore

$$\Lambda_\theta(\theta; u) = \alpha + B'(\theta) - C_\theta\big(m_{\text{req}}(u), \theta\big) \;\geq\; \alpha + B'(\theta) + \eta \;>\; 0,$$

using $B'(\theta) > 0$ from (P3). Thus $\theta \mapsto \Lambda(\theta; u)$ is strictly increasing on $[\underline{\theta}_0, u]$. If $\Lambda(\theta_1; u) < 0 < \Lambda(\theta_2; u)$ for some $\theta_1 < \theta_2$, continuity of $\Lambda(\cdot; u)$ in θ yields existence of a zero in (θ_1, θ_2), and strict monotonicity yields uniqueness. □

Lemma 5 (Monotonicity of the zero–crossing in the cut-off)**.** *Fix an interior belief cut-off* $u \in (\underline{\theta}_0, \overline{\theta}_0)$. *Let*

$$m^{\text{req}}(u) := \mu^{-1}(u), \qquad \Lambda(\theta; u) := \alpha\theta + B(\theta) - C\big(m^{\text{req}}(u), \theta\big).$$

Assume that for this u *the map* $\theta \mapsto \Lambda(\theta; u)$ *is strictly single-crossing, so the equation* $\Lambda(\theta; u) = 0$ *has a unique solution* $\theta^*(u) \in (\underline{\theta}_0, u)$. *Then* $u \mapsto \theta^*(u)$ *is continuously differentiable and*

$$\frac{d\theta^*}{du} = -\frac{\Lambda_u\big(\theta^*(u); u\big)}{\Lambda_\theta\big(\theta^*(u); u\big)} = \frac{\big(\eta + \kappa\Delta(u)\big)\, \mu^{-1\prime}(u)}{\alpha + B'(\theta^*(u)) + \eta + \kappa\Delta(u)} > 0\,,$$

where $\Delta(u) := m^{\text{req}}(u) - \theta^*(u) > 0$. *Hence the upper end of the pooling band* θ^* *moves strictly upward when the voting cut-off* u *is tightened.*

Proof. By definition of $\theta^*(u)$, we have the identity

$$\Lambda\big(\theta^*(u); u\big) = 0 \qquad \text{for all interior } u.$$

Because $u \in (\underline{\theta}_0, \overline{\theta}_0)$ and μ is strictly increasing and continuous, $m^{\text{req}}(u) = \mu^{-1}(u)$ is well defined and strictly increasing by (M1). Since $\theta^*(u) < u$, we have $\Delta(u) := m^{\text{req}}(u) - \theta^*(u) > 0$; i.e., the root lies on the *up-wedge*. By (S1)-(S4), C is C^1 off the diagonal $\{\Delta = 0\}$, so the partial derivatives below exist at $\big(\theta^*(u); u\big)$.

Differentiate the identity $\Lambda(\theta^*(u); u) = 0$ with respect to u. Using $\Lambda(\theta; u) = \alpha\theta + B(\theta) - C\big(\mu^{-1}(u), \theta\big)$ gives

$$\Lambda_u(\theta; u) = -\,C_m\big(\mu^{-1}(u), \theta\big)\, \mu^{-1\prime}(u), \qquad \Lambda_\theta(\theta; u) = \alpha + B'(\theta) - C_\theta\big(\mu^{-1}(u), \theta\big).$$

Evaluate at $\big(\theta^*(u); u\big)$. On the up-wedge ($\Delta > 0$), (S2)-(S4) imply

$$C_m\big(\mu^{-1}(u), \theta^*(u)\big) = \eta + \kappa\Delta(u) > 0, \qquad C_\theta\big(\mu^{-1}(u), \theta^*(u)\big) = -(\eta + \kappa\Delta(u)) < 0.$$

By (M1), $\mu^{-1\prime}(u) > 0$. Hence

$$\Lambda_u\big(\theta^*(u); u\big) = -\big(\eta + \kappa\Delta(u)\big)\,\mu^{-1\prime}(u) \ < \ 0,$$

and, using (P3) to get $B'(\theta) > 0$,

$$\Lambda_\theta\big(\theta^*(u); u\big) = \alpha + B'\big(\theta^*(u)\big) - C_\theta\big(\mu^{-1}(u), \theta^*(u)\big) = \alpha + B'\big(\theta^*(u)\big) + \eta + \kappa\Delta(u) \ > \ 0.$$

Since Λ is continuously differentiable in a neighborhood of the root and $\Lambda_\theta \neq 0$ there (strict single-crossing), the Implicit Function Theorem applies and yields

$$\frac{d\theta^*}{du} = -\frac{\Lambda_u\big(\theta^*(u); u\big)}{\Lambda_\theta\big(\theta^*(u); u\big)} = \frac{\big(\eta + \kappa\Delta(u)\big)\,\mu^{-1\prime}(u)}{\alpha + B'(\theta^*(u)) + \eta + \kappa\Delta(u)} \ > \ 0.$$

In particular, $0 < \frac{d\theta^*}{du} < \mu^{-1\prime}(u)$, with the bounds strict whenever $\Delta(u) > 0$ and $\eta > 0$. □

Lemma 6 (Pooling test). *Define the lower and upper endpoint tests:*

$$\Psi_\ell(u) \ := \ \Lambda\big(\underline{\theta}_0\,;\, u\big), \qquad \Psi_h(u) \ := \ \lim_{\theta\uparrow u} \Lambda(\theta; u) \ = \ \alpha\, u + B(u) - C\big(m^{\text{req}}(u), u\big).$$

Both are continuous in u; Ψ_ℓ is strictly decreasing whenever $m^{\text{req}}(u) > \underline{\theta}_0$ (Lemma 3). Moreover, $\Psi_\ell(u) < 0$ for every interior u with $\alpha\,\underline{\theta}_0 < 0$ and $B(\underline{\theta}_0) = 0$ (i.e. $b_0 = 0$; recall (P3) *normalizes $D(\underline{\theta}_0) = 0$). Then:*

$$\Pr(\text{partial pooling}) > 0 \quad \Longleftrightarrow \quad \Psi_\ell(u) < 0 \textit{ and } \Psi_h(u) > 0.$$

When $\Psi_h(u) \le 0$, the equilibrium is fully separating at the bar u (Proposition 2(iii)*).*

Proof. The lower endpoint satisfies $\Psi_\ell(u) \le \alpha\,\underline{\theta}_0 < 0$ (using $B(\underline{\theta}_0) = 0$, i.e. $b_0 = 0$, and $C \ge 0$), so $\Psi_\ell(u) < 0$ always holds for interior u.

If in addition $\Psi_h(u) > 0$, then $\Lambda(\cdot; u)$ is negative at $\underline{\theta}_0$ and positive as $\theta \uparrow u$. By the strict single-crossing property (Lemma 4), there exists a unique $\theta^*(u) \in (\underline{\theta}_0, u)$ with $\Lambda(\theta^*; u) = 0$, and the pooling mass $F(u) - F(\theta^*) > 0$.

If $\Psi_h(u) \le 0$, then $\Lambda(\theta; u) \le 0$ for all $\theta \le u$ (since Λ is strictly increasing in θ), so no type below u finds shading profitable and the equilibrium is fully separating. □

A.3 Additional Proofs

Corollary 4 (Degenerate voting rules). *Let $u = \underline{\theta}$ be the rule-specific belief cut-off.*

(i) ***Trivially easy rule*** $[u < \underline{\theta}_0]$. *Every posterior belief satisfies* $\mu(m) \geq u$*, hence minorities always vote "Yes". For the controller any message in the approval region is accepted; because* $C(m, \theta) \geq 0$ *with equality at* $m = \theta$*, truthful reporting maximises payoff but* any *message weakly below the truth is also a best reply. Consequently an uncountable continuum of Perfect-Sequential Equilibria exists, ranging from complete pooling* $m(\theta) \equiv m_L$ *(with any* $m_L \leq u$*) to full separation* $m(\theta) = \theta$*. The pooling probability can take any value in* $[0, 1]$.

(ii) ***Trivially tough rule*** $[u > \overline{\theta}_0]$. *No posterior belief can reach the hurdle, so minorities always vote "No". The controller's unique best reply is to choose the cheapest message, i.e.* $m(\theta) = \theta$ *(because* $C(m, \theta)$ *is minimised at* $m = \theta$*). The unique PSE is therefore fully separating with universal rejection.*

Lemma 7 (No single-message pooling)**.** *Fix an interior belief cut-off* $u \in (\underline{\theta}_0, \overline{\theta}_0)$*, and maintain Assumptions* P1-P3, S1-S4, M1. *Off the equilibrium path, beliefs satisfy the Grossman-Perry credible-beliefs refinement. Then there is* no *Perfect-Sequential Equilibrium (PSE) in which every* $\theta \in (\underline{\theta}_0, \overline{\theta}_0)$ *sends the same message* $m^\flat$*, regardless of whether minorities would approve or reject* $m^\flat$*.*

Proof. Let $m^{\text{req}} := \mu^{-1}(u)$ denote the minimal persuasive message (V3). Two cases:

(a) Common message is rejected: $\mu(m^\flat) < u$**.** Then minorities always vote "No" and the controller's equilibrium payoff is $U^{\text{eq}}(\theta) = -C(m^\flat, \theta)$. Pick a high type $\bar{\theta}$ close enough to $\overline{\theta}_0$ such that

$$\alpha\bar{\theta} + B(\bar{\theta}) \;>\; C\big(m^{\text{req}}, \bar{\theta}\big).$$

This is feasible because $\alpha\theta + B(\theta)$ is strictly increasing in θ (since $\alpha + B' > 0$ by (P3)), while the signaling cost $C(m^{\text{req}}, \bar{\theta})$ is bounded for fixed m^{req}. If type $\bar{\theta}$ deviates to the approved message m^{req} she obtains

$$U^{\text{dev}} = \alpha\bar{\theta} + B(\bar{\theta}) - C\big(m^{\text{req}}, \bar{\theta}\big) \;>\; 0 \;\geq\; U^{\text{eq}}(\bar{\theta}),$$

contradicting sequential rationality.

(b) Common message is approved: $\mu(m^\flat) \geq u$**.** We split by the location of $m^\flat$.

(b1) $m^\flat < \overline{\theta}_0$ *(some types under-report).* Choose any type $\theta^+ \in (m^\flat, \overline{\theta}_0]$ and fix a small $\varepsilon > 0$; consider the deviation $m^+ := m^\flat + \varepsilon$. Because the deviator is under-reporting at $m^\flat$, (S3) implies $C_m(m, \theta^+) = -\eta(\theta^+) < 0$ for all $m < \theta^+$, hence

$$C(m^+, \theta^+) < C(m^\flat, \theta^+).$$

Under the Grossman-Perry refinement, upon observing the *higher* off-path message m^+, minorities assign probability 1 to the set of types who benefit most from this deviation. Only *high* types with $\theta \geq m^\flat$ strictly benefit from moving up, so the posterior at m^+ is supported on $\{\theta \geq m^\flat\}$. Since the minority's approval statistic

$$g(\theta) := (1 - \alpha)\theta - B(\theta)$$

is strictly increasing in θ (P3), the posterior at m^+ first-order stochastically *dominates* the pooling posterior at $m^\flat$, and thus $\mathbb{E}[g(\theta) \mid m^+] \geq \mathbb{E}[g(\theta) \mid m^\flat]$. Because the project is approved at $m^\flat$, it remains approved at m^+ as well. The deviator's benefit is unchanged (approval), while her cost is strictly lower, so $U^{\text{dev}}(\theta^+) > U^{\text{eq}}(\theta^+)$–a contradiction.

(b2) $m^\flat \geq \overline{\theta}_0$ (all types over-report). Evaluate the pooling payoff at the bottom type. By normalization $B(\underline{\theta}_0) = 0$ (i.e. $b_0 = 0$) and since $m^\flat > \underline{\theta}_0$, (S2)-(S3) give $C(m^\flat, \underline{\theta}_0) > 0$. Because $\underline{\theta}_0 < 0$,

$$U^{\text{eq}}(\underline{\theta}_0) = \alpha\,\underline{\theta}_0 + B(\underline{\theta}_0) - C(m^\flat, \underline{\theta}_0) \;<\; 0.$$

This type strictly prefers truthful rejection (payoff 0), so a single-message approved pooling profile cannot be a PSE.

Combining (a), (b1), and (b2), no PSE can feature a single message used by all types when the cut-off u is interior. □

A.4 Credible-Beliefs Refinement (Grossman–Perry 1986)

A.4.1 Definition

Consider any signalling game with Perfect Bayesian Equilibria (PBE). Let m_{off} be a *message that is not sent with positive probability in the candidate equilibrium.* A PBE satisfies the credible-beliefs refinement if the receiver's belief $\mu(\cdot \mid m_{\text{off}})$ assigns probability 1 to the set

$$T^{\text{best}}(m_{\text{off}}) := \Big\{\theta \in \Theta \;:\; U_S\big(m_{\text{off}}, \theta\big) \;\geq\; U_S\big(m', \theta\big) \;\;\forall m' \text{ in the equilibrium support}\Big\},$$

i.e. the set of sender types that would (weakly) benefit *most* from deviating to m_{off}. If T^{best} is empty, beliefs are unrestricted.

Intuitively, the receiver should deem the unexpected signal most likely to come from the type with the strongest incentive to deviate, which is usually the *worst* type from the receiver's standpoint. This pessimism discourages profitable deviations and helps select a unique equilibrium.

A.4.2 Illustration in Our Partial-Pooling Equilibrium

Equilibrium recap. In the pooling band where $\theta < \theta^*$, all types send the common message m_L and minority shareholders reject the proposal. In the separating band where $\theta \geq \theta^*$, each type reports truthfully; the minorities approve precisely when $(1-\alpha)\theta - B(\theta) \geq \underline{\kappa}^{(\lambda)}$.

Off-path deviations. Consider a sender who instead transmits a higher message m_H that is not prescribed by the equilibrium. High-quality types ($\theta \geq \theta^*$) already secure approval without cost, so switching to m_H cannot improve their payoff. Low-quality types ($\theta < \theta^*$) are currently rejected and would gain $\alpha\theta + B(\theta) - C(m_H, \theta)$ only if the deviation were accepted. Hence the set of types with the strongest incentive to deviate, $T^{\text{best}}(m_H)$, contains exactly these low types.

Credible beliefs assign probability one to $\theta \in T^{\text{best}}(m_H)$ upon observing m_H. Minority shareholders therefore adopt the most pessimistic posterior, and the inequality $(1-\alpha)\mu(m_H)-\beta(m_H) \geq \underline{\kappa}^{(\lambda)}$ fails. They vote "No." The deviating sender then pays the full signalling cost $C(m_H, \theta)$ without gaining private benefits, making deviation strictly unprofitable.

Because every off-equilibrium message is treated in this way, no type wishes to deviate, and the partial-pooling profile remains the unique perfect sequential equilibrium under the given voting rule. Credible beliefs thus eliminate the need for arbitrary off-path assumptions: outsiders always interpret surprise signals in the manner most damaging to the sender, thereby ruling out equilibria that depend on optimistic beliefs and resolving multiplicity.

B Appendix B: Benchmark one-stage and bar-boundary results

B.1 Vanishing Exposure as $\lambda \to 0$

Under MoM, the voting test $(1-\alpha)\,\mu(m) - \beta^{\text{Div}}(m) \geq 0$ (cf. (V2)) does not involve λ, so ownership dispersion leaves the approval set unchanged. Under SM with quota $\pi \in (0,1)$, the rule constant in (V2) equals $\underline{\kappa}_{\text{SM}}^{(\lambda)} = -(1-2\pi)/\lambda$. Thus $\underline{\kappa}_{\text{SM}}^{(\lambda)} \to -\infty$ as $\lambda \downarrow 0$ when $\pi < \frac{1}{2}$ (lenient), and $\underline{\kappa}_{\text{SM}}^{(\lambda)} = (2\pi-1)/\lambda \to +\infty$ when $\pi > \frac{1}{2}$ (stringent). The next proposition formalizes these limits and provides explicit thresholds in λ.

Proposition 12 (Limit approval outcomes as $\lambda \to 0$)**.** *Work under* (P1)-(P3) *and* (M1). *Let* $\mu(m) = \mathrm{E}[\theta \mid m] \in [\underline{\theta}_0, \overline{\theta}_0]$ *and* $\beta^{\text{Div}}(m) = \mathrm{E}[D(\theta) \mid m] \in [D(\underline{\theta}_0), D(\overline{\theta}_0)]$. *Define the bounds*

$$\underline{L} \;:=\; (1-\alpha)\,\underline{\theta}_0 - D(\overline{\theta}_0), \qquad \overline{L} \;:=\; (1-\alpha)\,\overline{\theta}_0 - D(\underline{\theta}_0).$$

Then:

- *MoM ($\pi = \frac{1}{2}$). The approval inequality* (V2) *is independent of λ; hence the approval set (and any equilibrium allocation consistent with* (V2)*) is unchanged as $\lambda \to 0$.*
- *Lenient SM ($\pi < \frac{1}{2}$). Let*
$$\bar{\lambda}_{\text{len}} \;:=\; \frac{1-2\pi}{-\underline{L}} \;>\; 0.$$
For every $0 < \lambda \leq \bar{\lambda}_{\text{len}}$, the voting test (V2) *holds for every message m, so approval is certain:* $\Pr(\text{approval} \mid \lambda \leq \bar{\lambda}_{\text{len}}) = 1$.
- *Stringent SM ($\pi > \frac{1}{2}$). Let*
$$\bar{\lambda}_{\text{str}} \;:=\; \frac{2\pi-1}{\overline{L}} \;>\; 0.$$
For every $0 < \lambda \leq \bar{\lambda}_{\text{str}}$, the voting test (V2) *fails for every message m, so rejection is certain:* $\Pr(\text{approval} \mid \lambda \leq \bar{\lambda}_{\text{str}}) = 0$.

In particular, as $\lambda \downarrow 0$, lenient SM converges to universal approval, while stringent SM converges to universal rejection, whereas MoM is neutral in λ.

Proof. By (P1) and the definitions of posteriors, for every m,

$$\mu(m) \in [\underline{\theta}_0, \overline{\theta}_0], \qquad \beta^{\mathrm{Div}}(m) \in [D(\underline{\theta}_0),\, D(\overline{\theta}_0)].$$

Hence the left-hand side of (V2) is uniformly bounded:

$$(1-\alpha)\,\underline{\theta}_0 - D(\overline{\theta}_0) \;\le\; (1-\alpha)\,\mu(m) - \beta^{\mathrm{Div}}(m) \;\le\; (1-\alpha)\,\overline{\theta}_0 - D(\underline{\theta}_0). \tag{V4$'$}$$

MoM. For MoM, $\underline{\kappa}^{(\lambda)}_{\mathrm{MoM}} = 0$ in (V2), which does not depend on λ. Therefore the set $\{m : (1-\alpha)\mu(m) - \beta^{\mathrm{Div}}(m) \geq 0\}$ is independent of λ.

Lenient SM. For $\pi < \frac{1}{2}$, we have $\underline{\kappa}^{(\lambda)}_{\mathrm{SM}} = -(1-2\pi)/\lambda$. Using the lower bound in (V4$'$), the inequality (V2) holds for all m whenever

$$\underline{\kappa}^{(\lambda)}_{\mathrm{SM}} \le \underline{L} \quad\Longleftrightarrow\quad -\frac{1-2\pi}{\lambda} \le \underline{L}.$$

Since $\underline{L} < 0$ (because $\underline{\theta}_0 < 0$ and $D(\overline{\theta}_0) \geq 0$), this is equivalent to $\lambda \le (1-2\pi)/(-\underline{L}) = \bar{\lambda}_{\mathrm{len}}$. Thus for every $0 < \lambda \le \bar{\lambda}_{\mathrm{len}}$, (V2) is satisfied for all m, so approval is certain.

Stringent SM. For $\pi > \frac{1}{2}$, we have $\underline{\kappa}^{(\lambda)}_{\mathrm{SM}} = (2\pi-1)/\lambda$. Using the upper bound in (V4$'$), the inequality (V2) fails for all m whenever

$$\overline{L} \le \underline{\kappa}^{(\lambda)}_{\mathrm{SM}} \quad\Longleftrightarrow\quad \overline{L} \le \frac{2\pi-1}{\lambda}.$$

Because $\overline{L} > 0$ (as $\overline{\theta}_0 > 0$ and $D(\underline{\theta}_0) \geq 0$), this is equivalent to $\lambda \le (2\pi-1)/\overline{L} = \bar{\lambda}_{\mathrm{str}}$. (The bound $\overline{L}$ is never attained: achieving $(1-\alpha)\overline{\theta}_0$ requires the posterior to concentrate at $\overline{\theta}_0$, but then $\beta^{\mathrm{Div}} = D(\overline{\theta}_0) > D(\underline{\theta}_0)$, so the LHS of (V2) is strictly below $\overline{L}$.) Thus for every $0 < \lambda \le \bar{\lambda}_{\mathrm{str}}$, (V2) is violated for all m, so rejection is certain.

The three cases establish the claim. □

As ownership becomes extremely dispersed ($\lambda \to 0$), each minority investor's personal exposure to the project's payoff and to any diversion becomes negligible, so her private benefit from screening is too small to justify costly opposition–the classic free-rider logic in corporate control (Shleifer and Vishny (1986)). Under a lenient super-majority ($\pi < \frac{1}{2}$), the effective voting hurdle on minority support collapses in the atomistic limit, so even weak, lightly informed "Yes" support suffices and all proposals pass. Under a stringent super-majority ($\pi > \frac{1}{2}$), the hurdle explodes, and with atomistic owners unwilling to incur screening costs or coordinate, no proposal can clear the bar. By contrast, Majority-of-the-Minority fixes the belief cut-off independently of λ, placing it exactly at the knife-edge: greater dispersion neither weakens nor tightens the approval condition, so the approval set is unchanged as $\lambda \to 0$. However, as shown below, this knife–

edge neutrality does *not* carry over to the ordering of pooling regions: even holding dispersion fixed, different rules can generate systematically larger or smaller pooling masses once beliefs and signaling costs are taken into account.

Proposition 13 (Comparative statics of pooling mass at a *fixed* belief cut-off (transferable B; talk-up))**.** *Maintain* (P1)-(P3)*,* (S1)-(S4)*, and* (M1) *in the payoff-dominant (talk-up) regime* $0 < B'(\theta) < 1-\alpha$*. Fix a voting environment summarized by an interior belief cut-off* $u \le 0$ *(treated as* exogenous and fixed *throughout this proposition)*[11] *and let* $\theta^*(u) < 0$ *be the unique root of* $\Lambda(\theta; u) = 0$ *on* $(\underline{\theta}_0, u)$*, where*

$$\Lambda(\theta; u) \ := \ \alpha\,\theta + B(\theta) - C\big(m_{\text{req}}(u), \theta\big), \qquad m_{\text{req}}(u) = \mu^{-1}(u).$$

Define the pooling mass *(probability) as the mass of types who send* $m_{\text{req}}(u)$*:*

$$P(u) \ := \ F(u) \ - \ F\big(\theta^*(u)\big).$$

Consider the following payoff-side *parameter changes (holding the voting environment, and hence* u *and* $m_{\text{req}}(u)$*, fixed):*

- *(a) Upward quality shift: an additive shift* $\Delta > 0$ *in fundamentals so that the lying-profit becomes* $\Lambda_\Delta(\theta; u) := \alpha(\theta + \Delta) + B(\theta + \Delta) - C\big(m_{\text{req}}(u), \theta + \Delta\big)$*;*
- *(b) Higher ownership:* $\alpha \mapsto \alpha + \Delta\alpha$*, with* $\Delta\alpha > 0$*;*
- *(c) Upward* level *shift in private benefits:* $B(\theta) \mapsto B(\theta) + \Delta B$*, with* $\Delta B > 0$*.*

Then, at the fixed cut-off u*,*

$$\begin{aligned} \text{(a) Upward quality } (\Delta > 0): &\quad P \text{ strictly increases,} \\ \text{(b) Higher ownership } (\Delta\alpha > 0): &\quad P \text{ strictly decreases,} \\ \text{(c) Higher private benefits } (\Delta B > 0): &\quad P \text{ strictly increases.} \end{aligned}$$

Proof. All derivatives are computed at $\theta = \theta^*(u)$ with u fixed. By strict single-crossing on the

[11] If the belief cut-off u is *endogenous* to the parameter of interest via (V2)-(V3) (e.g., u depends on α or on the level of B through β), then the total derivative acquires an additional *bar-movement* term. Writing $P(u) = \int_{\theta^*(u)}^{u} f(\theta)\, d\theta$, the variable-limits Leibniz rule gives

$$P_u(u) = f(u) \ - \ f\big(\theta^*(u)\big)\,\theta^{*\prime}(u) \qquad \text{with} \qquad \theta^{*\prime}(u) = \frac{C_m\big(m_{\text{req}}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)}{\alpha + B'\big(\theta^*(u)\big) \ - \ C_\theta\big(m_{\text{req}}(u), \theta^*(u)\big)} \ > \ 0,$$

so for any parameter p one has

$$\frac{dP}{dp} = P_u(u)\,\frac{du}{dp} \ - \ f\big(\theta^*(u)\big)\,\frac{\partial\theta^*}{\partial p}\Big|_u,$$

where the second term is the *payoff-side* effect computed in the proposition (the "fixed-u" piece) and the first term is the *rule-side* effect through bar movement. Under additional structure (e.g., uniform prior with $\mu(m) = m$), $P_u(u) > 0$ along any pooling corridor, so the sign of dP/dp combines a positive bar-movement contribution with the payoff-side term; in general, the overall sign is parameter- and environment-dependent.

up-wedge (see Lemma 4), $\Lambda_\theta(\theta; u) > 0$ on $(\underline{\theta}_0, u]$; hence, by the Implicit Function Theorem,

$$\frac{d\theta^*}{d\,\text{param}} \;=\; -\,\frac{\Lambda_{\text{param}}}{\Lambda_\theta} \qquad \text{evaluated at } \theta = \theta^*(u).$$

Because $P(u) = F(u) - F\big(\theta^*(u)\big)$ with u fixed,

$$\frac{dP}{d\,\text{param}} \;=\; -\,f\big(\theta^*(u)\big)\,\frac{d\theta^*}{d\,\text{param}},$$

and by (P2) the density satisfies $f > 0$ on the support. We compute Λ_{param} case by case.

(a) Upward quality $\Delta > 0$. Here Δ enters *only* as a shift of the type argument in each payoff term, so by the chain rule

$$\Lambda_\Delta(\theta; u) = \alpha + B'(\theta + \Delta) - C_\theta\big(m_{\text{req}}(u), \theta + \Delta\big).$$

At $\Delta = 0$, $\Lambda_\Delta(\theta; u) = \alpha + B'(\theta) - C_\theta(m_{\text{req}}(u), \theta) = \Lambda_\theta(\theta; u)$. Thus, at the root,

$$\frac{d\theta^*}{d\Delta} = -\,\frac{\Lambda_\Delta}{\Lambda_\theta} = -1 \quad\Longrightarrow\quad \frac{dP}{d\Delta} = -f(\theta^*)\,\frac{d\theta^*}{d\Delta} = f(\theta^*) \;>\; 0.$$

(b) Ownership α. Differentiating $\Lambda(\theta; u) = \alpha\theta + B(\theta) - C(\cdot)$ with respect to α (holding u fixed) gives $\Lambda_\alpha(\theta; u) = \theta$. Evaluated at $\theta = \theta^*(u) < 0$ (since $u \le 0$ and the root lies below u in the talk-up regime), we have $\Lambda_\alpha < 0$, hence

$$\frac{d\theta^*}{d\alpha} = -\,\frac{\Lambda_\alpha}{\Lambda_\theta} > 0 \quad\Longrightarrow\quad \frac{dP}{d\alpha} = -f(\theta^*)\,\frac{d\theta^*}{d\alpha} \;<\; 0.$$

(c) Level shift in B. A constant increase $B(\theta) \mapsto B(\theta) + \Delta B$ adds one-for-one to Λ, i.e. $\Lambda_{\Delta B} = +1$. Therefore

$$\frac{d\theta^*}{d(\Delta B)} = -\,\frac{1}{\Lambda_\theta} < 0 \quad\Longrightarrow\quad \frac{dP}{d(\Delta B)} = -f(\theta^*)\,\frac{d\theta^*}{d(\Delta B)} \;>\; 0.$$

All inequalities are strict because $f(\theta^*) > 0$ by (P2) and $\Lambda_\theta > 0$ by strict single-crossing (Lemma 4). □

Hold the belief cut-off u fixed (so the "bar" minorities apply does not move). When projects are uniformly stronger ($\Delta > 0$), the approval prize $\alpha(\theta+\Delta)+B(\theta+\Delta)$ increases one-for-one at every type while the cost of the *same* minimal lie to $m_{\text{req}}(u)$ is unchanged in u; marginal types just below the bar now find it worthwhile to talk up. The indifferent type $\theta^*(u)$ therefore shifts *down*, widening $[\theta^*(u), u]$ and raising the pooling mass $P(u) = F(u) - F(\theta^*(u))$.

Raising ownership ($\alpha \uparrow$) makes the controller internalize more of the project payoff. Along the relevant (weakly negative) part of the support, this *reduces* the net prize from approval (because $\alpha\theta$ is more negative for $\theta < 0$), so fewer borderline types are willing to pay the signaling cost to clear the same bar; the indifferent type moves *up* and the pooling band shrinks, lowering $P(u)$.

A uniform upward shift in private benefits ($B \uparrow$) adds a constant to the prize conditional on approval without making the message any more persuasive under the fixed threshold. With a larger prize, more near-bar types choose the minimal persuasive message; $\theta^*(u)$ moves *down*, the pooling band widens, and $P(u)$ rises.

This subsection reports the one-stage fixed-enforcement benchmark and is included only to anchor the reversal result. The paper's central mechanism is the additional enforcement-feedback term under coupled rules $r(\pi)$. Detailed bar-boundary algebra, thin-tail conditions, GDWBR variants, and supplementary figures are collected in Appendix C.

B.2 Majority-of-Minority versus Supermajority Rules

The following propositions summarize the fixed-enforcement benchmark comparison between MoM and supermajority rules. We first show partial pooling under broad conditions when $\pi < 0.5$ and then characterize the stringent case $\pi > 0.5$, where sufficiently high quotas can induce full separation.

B.2.1 Partial Pooling and Separation under MoM and Supermajority Rules

Assumption 9 (Local strengthened CPC at the bar)**.** *Fix an interior belief cut-off $u \in (\underline{\theta}_0, \overline{\theta}_0)$ induced by the rule. Let $m^{\text{req}}(u) := \mu^{-1}(u)$. Assume*

$$\alpha\, u + B(u) \; > \; C\big(m^{\text{req}}(u), u\big). \tag{CPC$^+$}$$

Proposition 14 (Partial pooling under local CPC$^+$ for MoM and lenient SM)**.** *Maintain P1–P3, S1–3.2.3, and M1 in the payoff-dominant regime $0 < B'(\theta) < 1 - \alpha$, with $B(\underline{\theta}_0) = 0$. For a rule-specific interior belief cut-off $u \in (\underline{\theta}_0, \overline{\theta}_0)$ (either $u = u_{\text{MoM}}$ or $u = u_{\text{SM}}(\pi)$ with $\pi < \frac{1}{2}$), let $m^{\text{req}}(u) := \mu^{-1}(u)$ and define*

$$\Lambda(\theta; u) \; = \; \alpha\theta + B(\theta) - C\big(m^{\text{req}}(u), \theta\big).$$

Assume the local strengthened participation condition at the bar in Assumption CPC$^+$. Then for each rule (MoM or any lenient SM) there exists a unique $\theta^*(u) \in (\underline{\theta}_0, u)$ and the fixed-u PSE allocation *is partially pooling: truthful rejection on $[\underline{\theta}_0, \theta^*(u))$, pooling at $m^{\text{req}}(u)$ on $[\theta^*(u), u)$, and truthful approval on $[u, \overline{\theta}_0]$. Allocation uniqueness for a given u follows from Proposition 2*(ii).

Proof. We first prove that the lower endpoint of $\Lambda(\underline{\theta}_0; u)$ is negative. By (S2)-(S4) we have $C(\cdot, \theta) \geq 0$ with equality only at $m = \theta$, and by normalization $B(\underline{\theta}_0) = 0$. Hence

$$\Lambda(\underline{\theta}_0; u) = \alpha\, \underline{\theta}_0 + B(\underline{\theta}_0) - C\big(m^{\text{req}}(u), \underline{\theta}_0\big) \; \leq \; \alpha\, \underline{\theta}_0 \; < \; 0.$$

Next, we prove that the upper endpoint test $\Psi_h(u)$ is positive under local CPC$^+$. Define $\Psi_h(u) := \lim_{\theta \uparrow u} \Lambda(\theta; u) = \alpha u + B(u) - C\big(m^{\text{req}}(u), u\big)$. By Assumption CPC$^+$, $\Psi_h(u) > 0$.

By Lemma 4, the map $\theta \mapsto \Lambda(\theta; u)$ is continuous and strictly increasing on $[\underline{\theta}_0, u]$.

Given the upper and lower endpoints of $\Lambda(\underline{\theta}_0; u)$ and Lemma 4, by the intermediate value theorem, we have a unique zero $\theta^*(u) \in (\underline{\theta}_0, u)$ with $\Lambda(\theta^*(u); u) = 0$. Because $f > 0$ on the support (Assumption 1), the pooling mass $F(u) - F(\theta^*(u))$ is strictly positive. Proposition 2(i)-(iii) then yields the fixed-u PSE allocation: truthful rejection on $[\underline{\theta}_0, \theta^*(u))$, pooling at $m^{\text{req}}(u)$ on $[\theta^*(u), u)$, and truthful approval on $[u, \overline{\theta}_0]$. Allocation uniqueness for this u follows from Proposition 2(ii).
□

Intuitively, fix any interior belief cut–off u. If the local strengthened CPC at the bar ($\alpha u + B(u) > C\big(\mu^{-1}(u), u\big)$) holds, then the fixed-$u$ equilibrium exhibits *partial pooling*: some types just below u shade up to the minimal persuasive message and pass, while higher types separate truthfully. If this inequality fails–e.g., under a lenient SM that drives u sufficiently left so the approval gain shrinks while $C\big(\mu^{-1}(u), u\big) > 0$ (since $\mu^{-1}(u) \neq u$ in general)–the entire pooling region collapses and the fixed-u equilibrium is *separating* at that bar.

Proposition 15 (Stringent SM: partial pooling below a critical quota, full separation above)**.** *Maintain P1-P3, S1-S4, M1, and the tail-growth assumption on talk-up costs. Fix a* stringent *quota* $\pi \in (\frac{1}{2}, 1]$ *and let* $u_{\text{SM}}(\pi)$ *denote the induced (rule-specific) belief cut-off defined by* (V2)*; assume it is interior whenever invoked. Write* $m^{\text{req}}(u) := \mu^{-1}(u)$, $\Lambda(\theta; u) := \alpha\theta + B(\theta) - C\big(m^{\text{req}}(u), \theta\big)$, *and* $\theta^*(u) := \sup\{\theta < u : \Lambda(\theta; u) < 0\}$.

(i) Interior mapping. For any stringent π *with interior* $u_{\text{SM}}(\pi)$*, the interior cut-off is unique. Moreover, on its domain, the map* $\pi \mapsto u_{\text{SM}}(\pi)$ *is strictly increasing; if* $\beta \circ \mu^{-1}$ *is* C^1*, then* u_{SM} *is* C^1 *with*

$$\frac{du_{\text{SM}}}{d\pi} = \frac{2/\lambda}{(1-\alpha) - \beta_u\big(u_{\text{SM}}(\pi)\big)} > 0.$$

(ii) Dichotomy for stringent SM under CPC$^+$*. There exists* $\bar{\pi} \in (\frac{1}{2}, 1]$ *such that:*

$$\begin{cases} \textit{If } \frac{1}{2} < \pi < \bar{\pi},\ \text{CPC}^+ \textit{ holds at } u_{\text{SM}}(\pi) \implies \underline{\theta}_0 < \theta^*\big(u_{\text{SM}}(\pi)\big) < u_{\text{SM}}(\pi) \textit{ (partial pooling);} \\ \textit{If } \pi \geq \bar{\pi}, \implies \theta^*\big(u_{\text{SM}}(\pi)\big) = u_{\text{SM}}(\pi) \textit{ (full separation).} \end{cases}$$

If, in addition, $L(u) := u - \theta^*(u)$ *is strictly decreasing on the stringent corridor, then* $\bar{\pi}$ *is unique.*

Proof. (i) Interior mapping. By (V2), an interior cut-off u for a given π solves a strictly increasing equation in u (since $u \mapsto \beta(\mu^{-1}(u))$ is strictly increasing), which pins down uniqueness of $u_{\text{SM}}(\pi)$. Differentiating the (V2) identity yields the displayed derivative when $\beta \circ \mu^{-1}$ is C^1, and the denominator is positive by (P3); hence $\pi \mapsto u_{\text{SM}}(\pi)$ is strictly increasing.

(ii) Dichotomy. Fix a stringent π with interior $u := u_{\text{SM}}(\pi)$ and define

$$\Psi_\ell(u) := \Lambda(\underline{\theta}_0; u), \qquad \Psi_h(u) := \lim_{\theta \uparrow u} \Lambda(\theta; u) = \alpha u + B(u) - C\big(m^{\text{req}}(u), u\big).$$

Since $u > \underline{\theta}_0$ we have $m^{\mathrm{req}}(u) = \mu^{-1}(u) > \underline{\theta}_0$ and thus $C(m^{\mathrm{req}}(u), \underline{\theta}_0) > 0$ by (S2)-(S4); with $B(\underline{\theta}_0) = 0$ and $\alpha\underline{\theta}_0 < 0$, we get $\Psi_\ell(u) < 0$. Under CPC$^+$ at the induced bar, $\Psi_h(u) > 0$. By Lemma 4 (single crossing on $[\underline{\theta}_0, u]$), there is a unique $\theta^*(u) \in (\underline{\theta}_0, u)$, and Proposition 2 yields partial pooling for this u.

Next, consider tightening the quota. By (M1′) (full range, defined elsewhere) and tail growth of the gap costs, as $u \uparrow \overline{\theta}_0$ we have $m^{\mathrm{req}}(u) = \mu^{-1}(u) \uparrow +\infty$, so for $\theta \uparrow u$ the gap $\Delta(u,\theta) = m^{\mathrm{req}}(u) - \theta \to \infty$ and $C\big(m^{\mathrm{req}}(u), \theta\big) \to \infty$; hence $\Psi_h(u) \to -\infty$. By continuity of Ψ_h in u, there exists at least one interior $\bar{u}$ with $\Psi_h(\bar{u}) = 0$ (and $\Psi_h(u) > 0$ for $u < \bar{u}$ in a neighborhood, by the preceding argument).

Define $\bar{\pi}$ as the quota whose induced cut-off equals $\bar{u}$, i.e. $\bar{\pi}$ solves (V2) at $u = \bar{u}$. Part (i) implies $\pi \mapsto u_{\mathrm{SM}}(\pi)$ is strictly increasing; therefore: - If $\frac{1}{2} < \pi < \bar{\pi}$, then $u_{\mathrm{SM}}(\pi) < \bar{u}$ and $\Psi_h\big(u_{\mathrm{SM}}(\pi)\big) > 0$, yielding partial pooling. - If $\pi \geq \bar{\pi}$, then $u_{\mathrm{SM}}(\pi) \geq \bar{u}$ and $\Psi_h\big(u_{\mathrm{SM}}(\pi)\big) \leq 0$, so by Proposition 2(iii) the zero satisfies $\theta^*\big(u_{\mathrm{SM}}(\pi)\big) = u_{\mathrm{SM}}(\pi)$ (full separation).

If, moreover, $L(u) := u - \theta^*(u)$ is strictly decreasing along the stringent corridor, then $L(u_{\mathrm{SM}}(\pi))$ is strictly decreasing in π and the threshold is unique. □

The intuition behind Proposition 15 is as follows: Tightening the super-majority quota raises the belief cut-off u required for approval. To clear a higher bar, the controller must send a larger persuasive message $m_{\mathrm{req}}(u) = \mu^{-1}(u)$, which increases the lie $m_{\mathrm{req}}(u) - \theta$ for types just below u. With tail-growth in signaling costs, the expense of sufficiently large lies eventually overwhelms the bounded approval prize $\alpha u + B(u)$ (types reside on a compact support). At that point–corresponding to a critical quota $\bar{\pi}$–even the marginal type will not misreport, so the equilibrium becomes fully separating. For less strict quotas, CPC$^+$ ensures that the prize at the bar still exceeds the minimal signaling cost, so a nonempty set of near-threshold types shade up and partial pooling obtains.

B.2.2 Benchmark: Pooling under Uniform Prior and Identity Beliefs

As a benchmark, we establish the pooling ranking under the simplest information environment: uniform f and identity beliefs $\mu(m) = m$.

Proposition 16 (Uniform prior & identity beliefs: pooling rises with the bar; hence $P(u^{\mathrm{len}}_{\mathrm{SM}}) < P(u_{\mathrm{MoM}}) < P(u^{\mathrm{str}}_{\mathrm{SM}})$). *Maintain* (P1)-(P3), (S1)-(S4), *and* (M1) *in the payoff-dominant (up-talk) regime* $0 < B'(\theta) < 1 - \alpha$. *Assume the prior is uniform on* $[a, b]$ *with density* $f(\theta) = \frac{1}{b-a}$ *and identity beliefs* $\mu(m) = m$, *so the least-cost approved message at bar* u *is* $m^{\mathrm{req}}(u) = u$. *For each interior bar* $u \in (a, b)$ *define*

$$\Lambda(\theta; u) \;:=\; \alpha\,\theta + B(\theta) - C(u,\theta), \qquad \theta^*(u) \;:=\; \sup\{\theta < u \,:\, \Lambda(\theta; u) < 0\},$$

and the pooling mass

$$P(u) \;:=\; F(u) - F\big(\theta^*(u)\big) \;=\; \int_{\theta^*(u)}^{u} f(\theta)\,d\theta \;=\; \frac{u - \theta^*(u)}{b - a}.$$

Suppose the rule-induced cut-offs are interior and ordered as $u_{\mathrm{SM}}^{\mathrm{len}} < u_{\mathrm{MoM}} < u_{\mathrm{SM}}^{\mathrm{str}}$, and that partial pooling persists for every u on the two corridors $[u_{\mathrm{SM}}^{\mathrm{len}}, u_{\mathrm{MoM}}]$ and $[u_{\mathrm{MoM}}, u_{\mathrm{SM}}^{\mathrm{str}}]$ (so $\theta^*(u) \in (a, u)$ there). Then $P(u)$ is strictly increasing in u on each corridor. Consequently,

$$P\big(u_{\mathrm{SM}}^{\mathrm{len}}\big) \ < \ P\big(u_{\mathrm{MoM}}\big) \ < \ P\big(u_{\mathrm{SM}}^{\mathrm{str}}\big).$$

Proof. By Lemma 5, $0 < \theta^{*\prime}(u) < 1$ on any corridor with partial pooling (the denominator exceeds the numerator because $\alpha + B' > 0$). The key accounting identity is the variable-limits derivative

$$P'(u) \ = \ f(u) \ - \ f\big(\theta^*(u)\big)\, \theta^{*\prime}(u). \tag{5}$$

Under the uniform prior $f(u) = f(\theta^*(u)) = \frac{1}{b-a}$, so $P'(u) = \frac{1-\theta^{*\prime}(u)}{b-a} > 0$. Since $u_{\mathrm{SM}}^{\mathrm{len}} < u_{\mathrm{MoM}} < u_{\mathrm{SM}}^{\mathrm{str}}$ and P is strictly increasing along both corridors, the ranking follows. □

B.2.3 Economic intuition: threshold-bunching in shareholder approval votes

A shareholder vote creates a discontinuity: if the transaction fails, the controller gets no deal-related private benefits; if it passes, she obtains them. When disclosure and deal marketing can shift perceived quality, the controller's practical problem looks like clearing a *cliff*. In that environment, a wide range of "borderline" transactions naturally converge on the *minimum* persuasive story that is just sufficient to clear the pivotal shareholders' bar—because once the deal is barely over the line, making the story even rosier buys little additional support but can increase expected future costs.

In real-world proxy practice, this "minimum persuasive story" maps to a fairly standard approval package: a familiar bundle of banker analyses, management projections, comparables, and deal rationales calibrated to look *fair enough* to the voters who matter (Bhattarai et al., 2022). High-quality deals do not need to lean as hard on that playbook; very low-quality deals cannot profitably (or credibly) spin enough to get over the cliff. The strategic action happens in the middle, where many different underlying qualities can rationally choose the same "just-passing" disclosure package—so the observed messages look more alike even though fundamentals differ.

This is the same economics as *threshold-bunching* in corporate reporting and regulation: when rewards jump at a benchmark, agents often manage the observable metric to finish just above it, producing a pile-up right at the cutoff and "missing mass" just below (Burgstahler and Dichev, 1997; Degeorge et al., 1999).[12] A tougher voting rule moves the benchmark. Transactions that would have passed on straightforward disclosure become borderline and therefore have the strongest incentive to adopt the standardized "just-passing" package. In a one-stage benchmark with fixed enforcement costs, tightening the vote can therefore *increase* pooling: more deals are pushed into the "needs persuasion" region and they cluster at the minimum message that clears the vote.

[12] The general bunching methodology is developed in Saez (2010) and surveyed by Kleven (2016). We invoke the analogy only at the level of economic mechanism—a cliff incentive generates bunching—rather than importing the estimation toolkit.

This ranking can reverse under non-uniform priors or sluggish beliefs, as the formal analysis in Appendix C shows. But crucially, all of the above holds enforcement-cost parameters fixed ($r'(\pi) = 0$). The baseline puzzle is exactly why an ex post enforcement microfoundation matters: when the approval rule also changes the ease of settlement and ratification, the cost of "just-passing" persuasion rises with the vote threshold, and pooling can shrink even in the simplest information environment. The rest of the paper develops this enforcement/ratification channel.

B.2.4 Belief Sluggishness as an Alternative Channel (Robustness)

The GDWBR results (Proposition 18 in Appendix C) provide an alternative, complementary channel through which pooling can decrease with the bar: when minority beliefs are sufficiently sluggish (K small), the boundary responds aggressively to bar shifts even without the litigation channel. This mechanism operates through $(\mu^{-1})'(u) = 1/K$ in the numerator of $\theta^{*\prime}(u)$, and does not require r to covary with π. The two channels—litigation-adjusted enforcement and belief sluggishness—are additive: both contribute to making $\partial\theta^*/\partial u + (\partial\theta^*/\partial r) \cdot r'(\pi)/u'(\pi)$ larger. In environments with both sluggish beliefs *and* carry-forward ratification, LA-DWBR is easier to satisfy than either condition alone. Under the turnout microfoundation (Section **??**), small K corresponds to low attention/participation ($\tau < 1$), making GDWBR more empirically plausible.

C Bar–Boundary Tradeoff Analysis

This appendix collects the detailed bar–boundary comparison formulae, thin-tail sufficient conditions, GDWBR microfoundations, and illustrative figures that complement the benchmark result (Proposition 16) and the litigation-adjusted analysis (Proposition 3) in the main text.

C.1 Bar-Boundary Comparison Formula

As the sign of the $P'(u)$ depends on the crucially on the "bar-boundary" comparison outlined in Proposition 16 (Milgrom and Shannon (1994)), it is instructive to consider a more general form of the inequality:

$$\theta^{*\prime}(u) \lesseqgtr \frac{f(u)}{f(\theta^*(u))}$$

where

$$\theta^{*\prime}(u) = -\frac{\Lambda_u(\theta^*(u); u)}{\Lambda_\theta(\theta^*(u); u)}$$
$$= \frac{C_m\big(m^{\text{req}}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)}{\alpha + B'\big(\theta^*(u)\big) - C_\theta\big(m^{\text{req}}(u), \theta^*(u)\big)}$$
$$= \frac{C_m\big(\mu^{-1}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)}{\alpha + B'\big(\theta^*(u)\big) - C_\theta\big(\mu^{-1}(u), \theta^*(u)\big)}$$

from Lemma 5 represents the behavioral responsiveness of the marginal type to a small change in the bar u, and:

$$r(u) := \frac{f(u)}{f(\theta^*(u))}$$

is the density ratio at the two endpoints of the pooling band, capturing how the raw probability mass at the bar compares to the mass ”weight” at the moving boundary.

Both the terms $\frac{C_m\left(\mu^{-1}(u),\theta^*(u)\right)\,(\mu^{-1})'(u)}{\alpha+B'\left(\theta^*(u)\right)-C_\theta\left(\mu^{-1}(u),\theta^*(u)\right)}$ and $r(u)$ are economically meaningful. Insofar as the former term is concerned, consider the marginal type $\theta^*(u)$ who is just indifferent between telling the truth (rejected) and shading up to the bar u (approved). When the bar increases by $du > 0$, the required lie changes by $d\Delta = du - d\theta^*$. Holding the boundary fixed, this raises the marginal signaling burden by $C_m\big(\mu^{-1}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)\,du$ and therefore reduces the set of types willing to shade. The boundary then adjusts by $[\alpha + B'\big(\theta^*(u)\big) - C_\theta\big(\mu^{-1}(u), \theta^*(u)\big)]\,d\theta^*$ to restore indifference. This adjustment provides two offsetting benefits for higher types: (i) the truthful payoff increases at slope $\alpha + B'(\theta^*(u))$, and (ii) because gap costs satisfy $C_\theta = -C_m$, raising θ^* shrinks the lie and lowers the signaling cost at rate $C_m(\mu^{-1}(u), \theta^*(u))$. Rewriting $\theta^{*\prime}(u)$:

$$\theta^{*\prime}(u) = -\frac{\Lambda_u}{\Lambda_\theta} = \frac{C_m\big(\mu^{-1}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)}{\alpha + B'\big(\theta^*(u)\big) + C_m\big(\mu^{-1}(u), \theta^*(u)\big)}$$

The *numerator*, $C_m\big(\mu^{-1}(u), \theta^*(u)\big)\,(\mu^{-1})'(u)$, is the marginal cost pressure associated with the bar move du: it measures how much more expensive shading becomes per unit tightening. The *denominator*, $\alpha + B'\big(\theta^*(u)\big) + C_m\big(\mu^{-1}(u), \theta^*(u)\big)$, is the total marginal benefit associated with the boundary move $d\theta^*$: the first term $\alpha + B'$ is the truthful-payoff slope (higher types gain more from truth), and the $+C_m$ term captures the cost relief from reducing the lie for higher types (since $C_\theta = -C_m$ under gap costs).

The second term, $r(u) := f(u)/f(\theta^*(u))$, is a purely mechanical ”thickness” comparison of the prior at the two endpoints of the pooling band $[\theta^*(u), u]$. A small increase in the bar by du adds raw probability mass at the bar at rate $f(u)\,du$, while a same-sized movement of the lower endpoint would release mass at rate $f(\theta^*(u))\,du$. Thus $r(u)$ benchmarks how much mass the bar picks up per unit shift relative to what the boundary would release per unit shift, *holding behavior fixed.* Under a log–concave (unimodal) density with monotone shoulders, the benchmark depends on where the band sits: on the *right shoulder* (where $u > \theta^*(u)$ and f is weakly decreasing), we have $f(u) \leq f(\theta^*(u))$ so $r(u) \leq 1$, meaning the bar lies in a thinner tail and per–unit movements mechanically contribute less mass than at the boundary; on the *left shoulder* (where f is weakly increasing), $f(u) \geq f(\theta^*(u))$ so $r(u) \geq 1$, meaning the bar sits in a thicker region and per–unit movements mechanically contribute more mass than at the boundary. This density ratio is the benchmark in the bar–boundary comparison $P'(u) = f(u) - f(\theta^*(u))\, \theta^{*\prime}(u)$: it isolates the purely probabilistic (non-behavioral) side of how much mass a marginal shift of the bar can add relative to an equal shift of the band's lower endpoint.

C.2 Intuition behind Proposition 16

In the special case of Proposition 16 where $\mu(m) = m$, $\mu^{-1}(u) = u$, and $(\mu^{-1})'(u) = 1$. Thus, $\theta^{*\prime}(u)$ may be rewritten as $\theta^{*\prime}(u) = -\frac{\Lambda_u}{\Lambda_\theta} = \frac{C_m(u,\theta^*)}{\alpha + B'(\theta^*) + C_m(u,\theta^*)}$. Since $C_m(u, \theta^*) > 0$ and $\alpha + B'(\theta^*) > 0$ by derivation (Proposition 16), the denominator strictly exceeds the numerator: $0 < \theta^{*\prime}(u) < 1$. In words, the behavioral response moves the boundary by less than one-for-one when the bar shifts.

Turning to the mechanical side, under a uniform prior the density is flat, so raising the bar by du adds exactly one unit of mass per unit distance at the bar while an equal movement at the boundary releases mass at the same unit rate; equivalently, $r(u) = f(u)/f(\theta^*(u)) = 1$. Combining the two observations, $\theta^{*\prime}(u) < r(u)$: the boundary lags a one-for-one mass shift. Consequently, the pooling interval $[\theta^*(u), u]$ widens and the pooling mass $P(u)$ increases as we move from a lenient SM (lower u) to MoM (higher u), consistent with the bar–boundary comparison $P'(u) = f(u) - f(\theta^*(u))\, \theta^{*\prime}(u)$.

The same bar–boundary logic extends as we move from MoM to a stricter supermajority. Along any corridor where partial pooling persists, raising the bar u further continues to make the boundary respond less than one-for-one, so the pooling interval $[\theta^*(u), u]$ and hence the pooling mass $P(u)$ expand monotonically with u.

As shown in Figure 6, the bar advances faster than the boundary, so the gap $u - \theta^*(u)$ increases and the pooling mass $P(u)$ rises when moving from a lenient SM to MoM.

Figure 6: Pooling Probabilities and Lying-Profit Curves (Uniform f, Identity Beliefs)

(a) Panel A: Pooling Probabilities

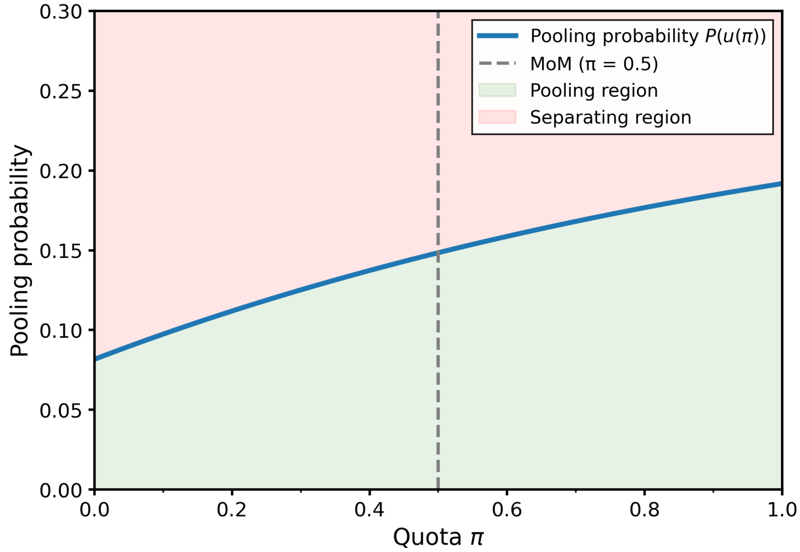

(b) Panel B: Lying-Profit Curves

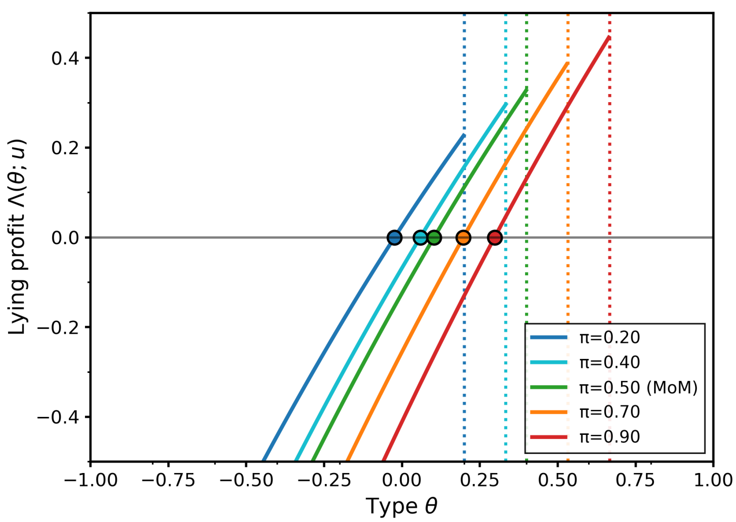

Panel A plots the simulated pooling probability $P(u(\pi))$ (solid blue) where the controller's cost is $C(m, \theta) = \eta|m - \theta| + \frac{\kappa}{2}(m - \theta)^2$; parameter values are $(\alpha, b_0, b_1, \lambda, K, \eta, \kappa) = (0.50, 0.12, 0.20, 0.10, 0.60, 0.20, 0.50)$. The grey dashed line marks the majority-of-the-minority threshold $\pi = 0.5$ and the green (pink) shading indicates the pooling (separating) region. Panel B shows the lying-profit function on the up-wedge, $\Lambda(\theta, u) = (\alpha + \beta)\theta - C(u)$, for $\pi \in \{0.20, 0.40, 0.50, 0.59, 0.61\}$; solid curves correspond to quotas that still permit pooling, dashed/dotted curves to quotas with complete separation, vertical dotted lines mark the induced belief cut-off $u = u_{\text{rule}}(\pi)$, and open markers indicate the type cutoff θ^* where $\Lambda(\theta, u) = 0$.

C.3 Log-Normal Densities and Thin Tails

We now turn to the comparative statics of the pooling region under single–peaked log–normal priors. With identity beliefs $\mu(m) = m$ and a log–concave density f whose mass is centered at the MoM bar (in our normalization $u_{\text{MoM}} = 0$), the idea is that most deals are middling while truly bad and truly great deals are rare (thin tails). A simple sufficient condition that formalizes "thin right tails"—and thus guarantees that the *boundary outruns the bar* once we tighten beyond MoM—is:

$$\sup_{u\in[0,\,u_{\text{SM}}^{\text{str}}]} \frac{f(u)}{f(\theta^*(u))} \le \underline{\gamma} \quad \text{where} \quad \underline{\gamma} := \frac{\eta}{\alpha + \overline{B'} + \eta} \in (0,1), \;\; \overline{B'} := \sup_{\theta} B'(\theta).$$

Intuitively, this "thin–tail" bound makes the density at the bar u uniformly small relative to the density at the moving boundary $\theta^*(u)$ on the right shoulder; because the marginal type's responsiveness satisfies $\theta^{*\prime}(u) \ge \underline{\gamma}$ (from gap costs and $0 < B'(\theta) < 1 - \alpha$), the density–weighted bar–boundary comparison $P'(u) = f(u) - f(\theta^*(u))\,\theta^{*\prime}(u)$ turns negative to the right of MoM. By contrast, on the left corridor (from a lenient SM up to MoM) the density is (weakly) increasing, so $f(u) \ge f(\theta^*(u))$ while $\theta^{*\prime}(u) < 1$, which makes $P'(u) > 0$. Putting the two shoulders together yields a clean ranking: a *lenient* supermajority generates the smallest pooling region; MoM induces strictly more pooling than any lenient SM; but any *strict* supermajority in turn induces strictly less pooling than MoM. In short, with sufficiently thin tails, MoM *maximizes* pooling among the three rules.

Proposition 17 (Up–talk, centered case: MoM maximizes pooling under thin right tails)**.** *Maintain* (P1)–(P3), (S1)–(S4)*, and* (M1) *in the payoff–dominant (up–talk) regime* $0 < B'(\theta) < 1 - \alpha$*. Assume identity beliefs* $\mu(m) = m$ *and affine private benefits* $B(\theta) = b_0 + b_1\theta$ *with* $b_1 \in (0, 1 - \alpha)$*. Work under the standing conventions of the paper that* $u_{\text{MoM}} = 0$ *and* 0 *is the unique mode of* f*, so* f *is (weakly) increasing on* $(-\infty, 0]$ *and (weakly) decreasing on* $[0, \infty)$*. Let* $u_{\text{SM}}^{\text{len}} < 0 < u_{\text{SM}}^{\text{str}}$ *be interior bars and suppose partial pooling persists on both corridors* $[u_{\text{SM}}^{\text{len}}, 0]$ *and* $[0, u_{\text{SM}}^{\text{str}}]$.

Define the lower–side pooling mass by $P(u) := \int_{\theta^*(u)}^{u} f(\theta)\,d\theta$*, where* $\theta^*(u) \in (\underline{\theta}_0, u)$ *is the unique boundary that solves* $\Lambda(\theta^*(u); u) = 0$*. Then:*

(a) ***Left corridor (**$u < 0$**).*** P *is strictly increasing as* $u \uparrow 0$*, i.e.*

$$P'(u) \;=\; f(u) \;-\; f(\theta^*(u))\,\theta^{*\prime}(u) \;>\; 0.$$

(b) ***Right corridor (**$u > 0$**).*** *If the right shoulder is* sufficiently thin *in the sense that*

$$\sup_{u\in[0,u_{\text{SM}}^{\text{str}}]} \frac{f(u)}{f(\theta^*(u))} \le \underline{\gamma} \quad \textit{with} \quad \underline{\gamma} := \frac{\eta}{\alpha + \overline{B'} + \eta} \in (0,1), \;\; \overline{B'} := \sup_{\theta} B'(\theta),$$

then P *is (weakly) decreasing as* u *moves right from* 0*, i.e.*

$$P'(u) \;=\; f(u) \;-\; f(\theta^*(u))\,\theta^{*\prime}(u) \;\le\; 0.$$

Consequently, on any pair of pooling corridors to the left and right of 0*, the pooling mass attains its maximum at the MoM bar:*

$$P(u) \text{ is maximized at } u = 0 \; (= u_{\text{MoM}}).$$

Proof. By the variable–limits rule,

$$P'(u) \;=\; f(u) \;-\; f(\theta^*(u))\,\theta^{*\prime}(u).$$

Under identity beliefs ($m^{\text{req}}(u) = u$) and gap costs (S4), Lemma 5 gives

$$\theta^{*\prime}(u) \;=\; \frac{C_m(u, \theta^*(u))}{\alpha + B'(\theta^*(u)) + C_m(u, \theta^*(u))} \;\in\; (0, 1).$$

Since on the up–wedge $C_m(u, \theta^*) = \eta + \kappa\Delta(u) \geq \eta$ with $\Delta(u) := u - \theta^*(u) > 0$ and $B'(\theta^*) \leq \overline{B'}$, we obtain the uniform lower bound

$$\theta^{*\prime}(u) \;\geq\; \underline{\gamma} \;:=\; \frac{\eta}{\alpha + \overline{B'} + \eta} \in (0, 1).$$

(a) Left corridor. For $u < 0$, the band $[\theta^*(u), u]$ sits on the increasing shoulder, so $f(u) \geq f(\theta^*(u))$. With $0 < \theta^{*\prime}(u) < 1$,

$$P'(u) \;=\; f(u) - f(\theta^*(u))\,\theta^{*\prime}(u) \;\geq\; f(\theta^*(u))\big(1 - \theta^{*\prime}(u)\big) \;>\; 0.$$

(b) Right corridor. For $u > 0$, write $r(u) := f(u)/f(\theta^*(u)) \in (0, 1]$. Then

$$P'(u) \;=\; f(\theta^*(u))\big[r(u) - \theta^{*\prime}(u)\big] \;\leq\; f(\theta^*(u))\big[\underline{\gamma} - \theta^{*\prime}(u)\big] \;\leq\; 0,$$

whenever $r(u) \leq \underline{\gamma}$. Thus P (weakly) decreases on $[0, u^{\text{str}}_{\text{SM}}]$ under the stated thin–tail condition.

Combining (a) and (b), P rises up to $u = 0$ and (weakly) falls thereafter, so it is maximized at $u = 0 = u_{\text{MoM}}$. □

A MoM threshold placed at the center of the distribution induces the largest pool of "borderline" controllers to shade up: the lying band sits exactly in the thickest part of the quality distribution and the lie to reach the bar is still affordable. A more lenient bar lets many of those mediocre deals pass truthfully, and a more stringent bar makes the required lie so costly that only a handful of very good deals in the thin right tail are still willing to pool. Thus, ironically, MoM can be the noisiest rule in terms of pooling.

As shown in Figure 7, the boundary advances faster than the bar on the right corridor, so pooling decreases when we move from the MoM regime to the strict SM regime.

Figure 7: Pooling Probabilities and Lying-Profit Curves (Logconcave, Identity Beliefs)

(a) Panel A: Pooling Probabilities

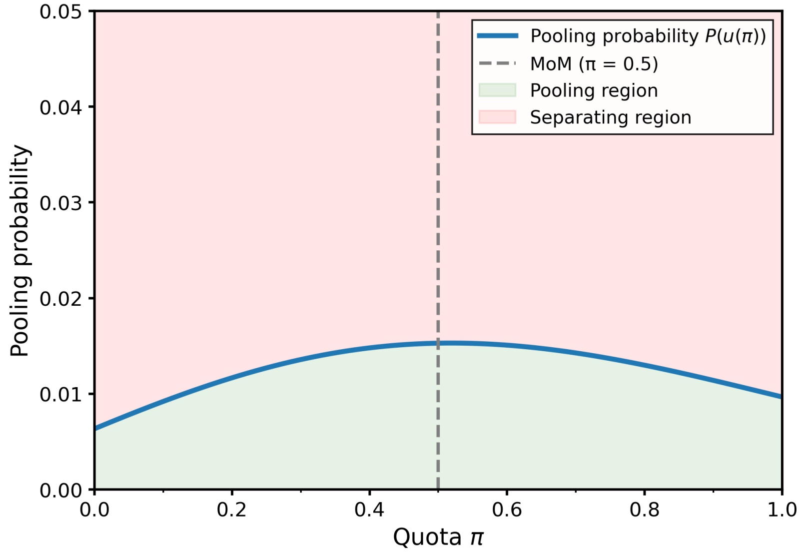

(b) Panel B: Lying-Profit Curves

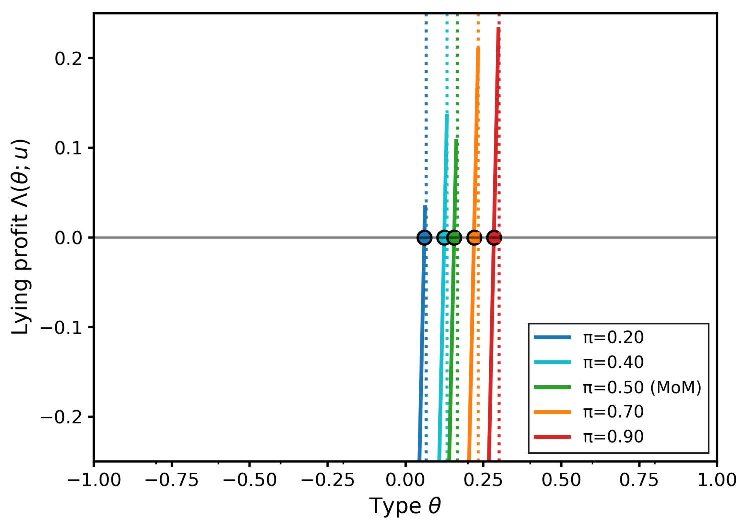

Panel A plots the simulated pooling probability $P(u(\pi))$ (solid blue) where the controller's cost is $C(m,\theta) = \eta|m - \theta| + \frac{\kappa}{2}(m-\theta)^2$; parameter values are $(\alpha, b_0, b_1, \lambda, K, \eta, \kappa) = (0.50, 0.12, 0.20, 0.20, 0.60, 20, 0.10)$. The density function is $f(\theta) = \frac{1}{0.20\sqrt{2\pi}\,[\Phi(1/0.20) - \Phi(-1/0.20)]} \exp\big(-\theta^2/(2\cdot 0.20^2)\big)\, \mathbf{1}_{[-1,1]}(\theta)$. The grey dashed line marks the majority-of-the-minority threshold $\pi = 0.5$ and the green (pink) shading indicates the pooling (separating) region. Panel B shows the lying-profit function on the up-wedge, $\Lambda(\theta, u) = (\alpha+\beta)\theta - C(u)$, for $\pi \in \{0.20, 0.40, 0.50, 0.59, 0.61\}$; solid curves correspond to quotas that still permit pooling, dashed/dotted curves to quotas with complete separation, vertical dotted lines mark the induced belief cut-off $u = u_{\text{rule}}(\pi)$, and open markers indicate the type cutoff θ^* where $\Lambda(\theta, u) = 0$.

C.4 GDWBR Microfoundations and Main Result

The results in this subsection hold the enforcement-cost parameters fixed (equivalently, $r'(\pi) = 0$). They provide one channel—belief sluggishness—through which the bar–boundary comparison can reverse. The main text (Section 5.4) introduces the complementary litigation channel that operates through $r'(\pi) > 0$.

The relative sizes of $r(u)$ and $\theta^{*\prime}(u)$ depend on the primitives—namely the shape of the prior f, the belief map μ, and the marginal costs and benefits of shading to the bar u (see Section C.1). To obtain clean, tractable comparative statics across the pooling regions induced by different voting rules, we provide microfounded sufficient conditions under which the bar–boundary comparison is signable. The key additional feature we assume here is that of *sluggish beliefs*, captured by a local responsiveness parameter $\mu'(m_{\text{req}}(u)) = K \in (0, 1]$, so that minorities do not update one–for–one with shifts in the belief cut–off. This behavioral feature can be motivated by standard theories of noisy signaling, rational inattention, or endogenous non–participation, and it enters the boundary–slope expression multiplicatively through $(\mu^{-1})'(u) = 1/K$, delivering a transparent dial that links primitives to the density–weighted responsiveness test (Blume et al., 2007; Sims, 2003; Matějka and Tabellini, 2021; Palfrey and Rosenthal, 1983; Osborne et al., 2000; Levine and Palfrey, 2007).

Under the turnout microfoundation in Section **??**, local responsiveness at the bar satisfies

$$\mu'(m_{\text{req}}(u)) = \tau\, \bar{\mu}'(m_{\text{req}}(u)).$$

In the identity-beliefs benchmark where $\bar{\mu}(m) = m$ (and $\mu_0 = 0$), the effective posterior is $\mu(m) = \tau m$, so $K = \tau$. More generally, K can be interpreted as the effective attention-weighted slope at the bar, combining turnout with the fully attentive posterior slope. For closed-form derivations we retain the linear reduced form $\mu(m) = Km$; this linear specification can be viewed as the local linearization induced by turnout around the bar.

We assume $\mu(m) = K\, m$ with $K \in (0, 1]$, the gap–cost specification from (S4), and affine private benefits $B(\theta) = b_0 + b_1\theta$. Under these primitives, $(\mu^{-1})'(u) = 1/K$. Substituting into the general boundary–slope formula from Lemma 5 yields the explicit expression

$$\theta^{*\prime}(u) = \frac{\eta + \kappa\,\Delta(u)}{\alpha + b_1 + \eta + \kappa\,\Delta(u)} \cdot \frac{1}{K}, \qquad K \in (0, 1], \quad \Delta(u) > 0.$$

Hence $\theta^{*\prime}(u) \in \big(0,\ 1/K\big)$, it is strictly increasing in the required lie $\Delta(u)$, and

$$\lim_{\Delta \downarrow 0} \theta^{*\prime}(u) = \frac{\eta}{\alpha + b_1 + \eta} \cdot \frac{1}{K}.$$

Under log–concavity with a unique mode at u_{MoM}, the shoulder ratio $r(u) = f(u)/f(\theta^*(u))$ satisfies $r(u) \geq 1$ on the left corridor $u \in [u_{\text{SM}}(\pi), u_{\text{MoM}}]$ and $r(u) \leq 1$ on the right corridor $u \in [u_{\text{MoM}}, \bar{u}]$. Because $\theta^{*\prime}(u)$ is increasing in the required lie $\Delta(u)$, its *minimum* on any corridor occurs as $\Delta(u) \downarrow 0$. Hence a simple global bound at the corridor's thinnest lie delivers a *sufficient condition* for the density–weighted comparison:

$$\frac{\eta}{K\big(\alpha + b_1 + \eta\big)} \;\geq\; \sup_{u \text{ on the comparison path}} r(u)$$

Under this bound we have $\theta^{*\prime}(u) \geq r(u)$ for all interior u on the path, i.e. the generalized density–weighted boundary responsiveness Assumption 10 holds.[13]

Assumption 10 (Global density–weighted boundary responsiveness (GDWBR))**.** *In the payoff–dominant (talk–up) regime, for every interior bar u at which partial pooling holds and $\theta^*(u) \in (\underline{\theta}_0, u)$, impose the global inequality*

$$\theta^{*\prime}(u) \;\geq\; \frac{f(u)}{f(\theta^*(u))}$$

for all u along the lenient corridor $u \in [u_{\mathrm{SM}}(\pi), u_{\mathrm{MoM}}]$ and the strict corridor $u \in [u_{\mathrm{MoM}}, \bar{u}]$, while pooling persists.

Intuitively, when beliefs are sluggish ($\mu'(m_{\mathrm{req}}(u)) = K \in (0, 1]$) the boundary slope picks up the multiplier $(\mu^{-1})'(u) = 1/K$, so a given change in the bar requires a larger signal adjustment to restore indifference. Economically, slower updating by minorities makes near–threshold senders shade more aggressively, which pushes the marginal type (the boundary) to move faster. As K falls, the left–hand side $\eta/[K(\alpha + b_1 + \eta)]$ rises, making it easier to dominate the density ratio $\sup_{u \text{ on path}} r(u)$ and thus satisfy the global condition.

Proposition 18 (Global DWBR: pooling strictly decreases from lenient SM → MoM → strict SM)**.** *Maintain* (P1)–(P3), (S1)–(S4), *and* (M1) *in the payoff–dominant (up–talk) regime* $0 < B'(\theta) < 1-\alpha$. *For each interior belief cut–off $u \in (\underline{\theta}_0, \overline{\theta}_0)$ let the least–cost approved message be $m^{\mathrm{req}}(u) := \mu^{-1}(u)$ and define the lying–profit $\Lambda(\theta; u) := \alpha\,\theta + B(\theta) - C\big(m^{\mathrm{req}}(u), \theta\big)$ and the unique lower boundary $\theta^*(u) \in (\underline{\theta}_0, u)$ by $\Lambda\big(\theta^*(u); u\big) = 0$ (Lemma 4). Let the pooling mass below the bar be*

$$P(u) := \int_{\theta^*(u)}^{u} f(\theta)\, d\theta.$$

Assume the rule–induced cut–offs are interior and ordered as $u^{\mathrm{len}}_{\mathrm{SM}} < u_{\mathrm{MoM}} < u^{\mathrm{str}}_{\mathrm{SM}}$, and that partial pooling persists on both corridors $[u^{\mathrm{len}}_{\mathrm{SM}}, u_{\mathrm{MoM}}]$ and $[u_{\mathrm{MoM}}, u^{\mathrm{str}}_{\mathrm{SM}}]$ (so $\theta^*(u) \in (\underline{\theta}_0, u)$ there). Suppose the Global Density–Weighted Boundary Responsiveness *(GDWBR) condition holds on* both *corridors:*

$$\theta^{*\prime}(u) \;\geq\; \frac{f(u)}{f\big(\theta^*(u)\big)} \qquad \textit{for every interior } u \textit{ on the two corridors.}$$

Then $P(u)$ is strictly decreasing in u on each corridor, and therefore

$$P\big(u^{\mathrm{len}}_{\mathrm{SM}}\big) \;>\; P\big(u_{\mathrm{MoM}}\big) \;>\; P\big(u^{\mathrm{str}}_{\mathrm{SM}}\big).$$

[13]For a uniform prior, $r(u) = 1$, so the sufficient condition reduces to $\theta^{*\prime}(u) \geq 1$ along the path.

Each inequality is weak-to-strong: it is strict whenever GDWBR holds with strict inequality on a positive-measure subset of the relevant corridor (which obtains generically when f is non-constant).

Proof. By single–crossing on the up–wedge (Lemma 4), the boundary $\theta^*(u) \in (\underline{\theta}_0, u)$ is unique and differentiable in u along any pooling corridor (Implicit Function Theorem applied to $\Lambda(\theta; u) = 0$ with $\Lambda_\theta > 0$). Differentiate $P(u) = \int_{\theta^*(u)}^{u} f(\theta)\, d\theta$ using the variable–limits Leibniz rule:

$$P'(u) = f(u) - f\big(\theta^*(u)\big)\, \theta^{*\prime}(u).$$

By GDWBR, $\theta^{*\prime}(u) \geq f(u)/f(\theta^*(u))$ on each corridor, implying

$$P'(u) \;\leq\; f(u) - f\big(\theta^*(u)\big)\, \frac{f(u)}{f(\theta^*(u))} \;=\; 0,$$

with strict inequality whenever GDWBR holds strictly on a positive-measure subset of the corridor. Hence P is (strictly) decreasing in u along $[u_{\text{SM}}^{\text{len}}, u_{\text{MoM}}]$ and along $[u_{\text{MoM}}, u_{\text{SM}}^{\text{str}}]$. Because $u_{\text{SM}}^{\text{len}} < u_{\text{MoM}} < u_{\text{SM}}^{\text{str}}$, integration of $P'(u) \leq 0$ over the two corridors yields

$$P\big(u_{\text{SM}}^{\text{len}}\big) \;>\; P\big(u_{\text{MoM}}\big) \;>\; P\big(u_{\text{SM}}^{\text{str}}\big),$$

with the stated strictness condition. □

C.5 GDWBR Intuition and Illustration

Tightening the belief cut–off u (which is monotone in the voting rule) shifts the marginal type $\theta^*(u)$ who is just indifferent between telling the truth (and being rejected) and shading up to u (and being approved). With sufficiently sluggish beliefs, minority investors underreact to any given signal, so a higher bar requires a disproportionately stronger persuasive message. For types just below the bar, the marginal benefit from shading then rises faster than the marginal cost, making the boundary respond strongly: the derivative $\theta^{*\prime}(u)$ becomes large.

Assumption 10 formalises this by requiring that the behavioral response $\theta^{*\prime}(u)$ dominates the purely mechanical mass effect at the bar in the density–weighted sense, $\theta^{*\prime}(u) \geq f(u)/f(\theta^*(u))$. In the bar–boundary accounting $P'(u) = f(u) - f\big(\theta^*(u)\big)\, \theta^{*\prime}(u)$, this dominance implies $P'(u) \leq 0$: as the rule tightens (higher u), the lower endpoint of the pooling band chases the bar quickly enough that the interval $[\theta^*(u), u]$ shrinks and pooling mass falls. Along any corridor where partial pooling persists, a lenient SM (low u) therefore sustains more pooling than MoM, and MoM in turn sustains more pooling than a stricter SM (high u).

Figure 8 illustrates the GDWBR case: the boundary's faster movement closes the gap $u - \theta^*(u)$, so pooling decreases when tightening from lenient SM to MoM and from MoM to strict SM, until the band vanishes.

Figure 8: Pooling Probabilities and Lying-Profit Curves (Global DWBR)

(a) Panel A: Pooling Probabilities

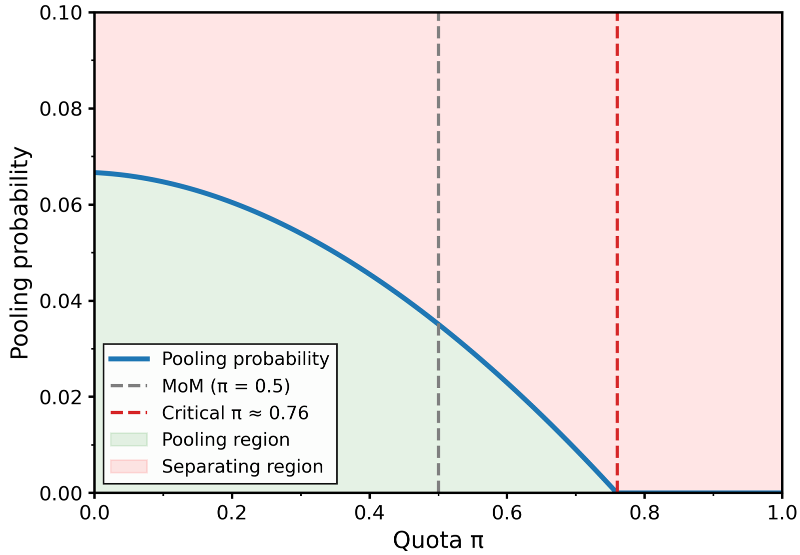

(b) Panel B: Lying-Profit Curves

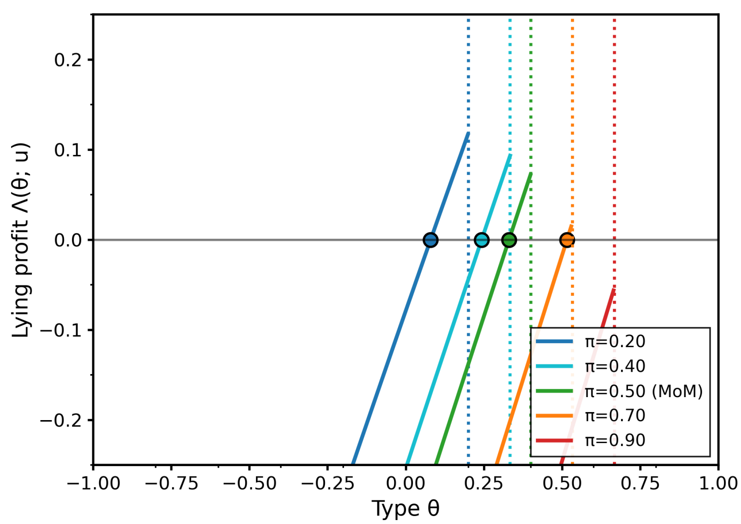

Panel A plots the simulated pooling probability $P(u(\pi))$ (solid blue) under a uniform type distribution on $[\theta_0, \theta_1] = [-1, 1]$ with sluggish beliefs $\mu(m) = Km$, where Assumption 10 holds. The controller's cost is $C(m, \theta) = \eta|m - \theta| + \frac{\kappa}{2}(m - \theta)^2$; parameter values are $(\alpha, b_0, b_1, \lambda, K, \eta, \kappa) = (0.50, 0.12, 0.20, 0.10, 0.25, 0.20, 0.12)$. The grey dashed line marks the majority-of-the-minority threshold $\pi = 0.5$, the red dashed line marks the critical quota $\pi \approx 0.76$ at which $P(u(\pi))$ hits zero, and the green (pink) shading indicates the pooling (separating) region. Panel B shows the lying-profit function on the up-wedge, $\Lambda(\theta, u) = (\alpha + \beta)\theta - C(u)$, for $\pi \in \{0.20, 0.40, 0.50, 0.70, 0.90\}$; vertical dotted lines mark the induced belief cut-off $u = u_{\text{rule}}(\pi)$, and open markers indicate the type cut-off θ^* where $\Lambda(\theta^*, u) = 0$.

D Appendix C: Formal proofs for private-ordering and committee extensions

D.1 Private-ordering proofs

D.1.1 Proof of Proposition 7

Proof. By definition,

$$V_C(H;\alpha,v) - V_C(L;\alpha,v) = \alpha\big(V_M(H) - V_M(L)\big) + \big(\Pi(H;v) - \Pi(L;v)\big) = \alpha\Delta_M + \Delta_\Pi(v).$$

Assumption 7 gives $\Delta_M > 0$. Hence $\alpha \mapsto V_C(H;\alpha,v) - V_C(L;\alpha,v)$ is affine and strictly increasing. The unique zero is

$$\alpha^*(v) := -\frac{\Delta_\Pi(v)}{\Delta_M} > 0,$$

because $\Delta_\Pi(v) < 0$. Therefore

$$V_C(H;\alpha,v) \geq V_C(L;\alpha,v) \iff \alpha \geq \alpha^*(v).$$

So the controller chooses H iff $\alpha \geq \alpha^*(v)$, proving monotone comparative statics in α. □

D.1.2 Proof of Proposition 8

Proof. From Proposition 7,

$$\alpha^*(v) = -\frac{\Delta_\Pi(v)}{\Delta_M}, \qquad \Delta_M > 0.$$

Take $v_2 > v_1$. Assumption 7 implies $\Delta_\Pi(v_2) \leq \Delta_\Pi(v_1)$. Multiplying by -1 and dividing by positive Δ_M yields

$$\alpha^*(v_2) \geq \alpha^*(v_1).$$

Hence $\alpha^*(v)$ is weakly increasing in v. Equivalently, fixing α, the set

$$\{v : \alpha < \alpha^*(v)\}$$

is increasing in v, so larger private-benefit intensity weakly expands the region where L is chosen. □

D.1.3 Proof of Proposition 9

Proof. Assumption 8 gives $W(H) > W(L)$, so welfare chooses H. If $\alpha < \alpha^*(v)$, Proposition 7 implies $V_C(L;\alpha,v) > V_C(H;\alpha,v)$, so the controller chooses L. Therefore controller choice and welfare choice differ:

$$\arg\max_{G\in\{L,H\}} V_C(G;\alpha,\nu) \neq \arg\max_{G\in\{L,H\}} W(G).$$

□

D.2 Committee monotonicity results

The following seven results are referenced in Section 6.6.3.

Lemma 8 (Gate-induced MLR and FOSD order)**.** *For $0 \le \iota_1 < \iota_2 \le 1$, the likelihood ratio $f_{\iota_2}(\theta)/f_{\iota_1}(\theta)$ is increasing in θ. Therefore f_{ι_2} first-order stochastically dominates f_{ι_1}.*

Proof. For $j \in \{1,2\}$, write

$$q_j(\theta) := q_{\iota_j}(\theta) = 1 - \iota_j + \iota_j Q(\theta), \qquad f_j(\theta) := f_{\iota_j}(\theta) = \frac{q_j(\theta) f(\theta)}{Z_j},$$

where $Z_j := \int q_j(t) f(t)\,dt > 0$. Then

$$\frac{f_2(\theta)}{f_1(\theta)} = \frac{Z_1}{Z_2} \cdot \frac{q_2(\theta)}{q_1(\theta)}.$$

The constant Z_1/Z_2 does not depend on θ, so monotonicity is determined by

$$h(\theta) := \frac{q_2(\theta)}{q_1(\theta)} = \frac{1-\iota_2+\iota_2 Q(\theta)}{1-\iota_1+\iota_1 Q(\theta)}.$$

Differentiate with respect to Q:

$$\frac{d}{dQ}\left(\frac{1-\iota_2+\iota_2 Q}{1-\iota_1+\iota_1 Q}\right) = \frac{\iota_2-\iota_1}{(1-\iota_1+\iota_1 Q)^2} > 0$$

since $\iota_2 > \iota_1$. Because $Q'(\theta) > 0$, $h(\theta)$ is increasing in θ, so f_2/f_1 is increasing in θ. This is MLR order. MLR implies first-order stochastic dominance, so f_{ι_2} FOSD-dominates f_{ι_1}. □

Proposition 19 (Gate monotonicity for low-type mass)**.** *For every cutoff $c \in (\underline{\theta}_0, \overline{\theta}_0)$,*

$$F_\iota(c) := \Pr_\iota(\theta \le c) = \int_{\underline{\theta}_0}^{c} f_\iota(\theta)\,d\theta$$

is weakly decreasing in ι. In particular, for any fixed voting bar u, the needs-persuasion mass $\Pr_\iota(\theta < u)$ is weakly decreasing in committee independence.

Proof. By Lemma 8, for $\iota_2 > \iota_1$, f_{ι_2} first-order stochastically dominates f_{ι_1}. By the definition of

FOSD, for every cutoff c:

$$F_{\iota_2}(c) = \Pr_{\iota_2}(\theta \le c) \le \Pr_{\iota_1}(\theta \le c) = F_{\iota_1}(c).$$

Hence $F_\iota(c)$ is weakly decreasing in ι. Setting $c = u$ gives $\Pr_\iota(\theta < u)$ weakly decreasing in ι. □

Proposition 20 (Gate channel and expected enforcement burden at fixed bar)**.** *Fix (u, r, σ) and define*

$$\Psi_u(\theta) := \mathbf{1}\{\theta < u\} L(u - \theta, r, \sigma).$$

If $L_g(g, r, \sigma) \ge 0$ on $g \ge 0$, then $\mathrm{E}_\iota[\Psi_u(\theta)]$ is weakly decreasing in ι.

Proof. Fix (u, r, σ). For $\theta < u$, $\Psi_u(\theta) = L(u - \theta, r, \sigma)$ and

$$\frac{d}{d\theta}\Psi_u(\theta) = -L_g(u - \theta, r, \sigma) \le 0$$

by $L_g \ge 0$. For $\theta \ge u$, $\Psi_u(\theta) = 0$. So Ψ_u is weakly decreasing on $[\underline{\theta}_0, \overline{\theta}_0]$.

From Lemma 8, larger ι yields FOSD shifts upward in θ. Under FOSD, expectation of any weakly decreasing function is weakly lower. Therefore $\mathrm{E}_\iota[\Psi_u(\theta)]$ is weakly decreasing in ι. □

Lemma 9 (Signal recommendation is monotone informative in θ)**.** *For fixed $\iota > 0$, $p_\iota(\theta)$ is increasing in θ. Therefore the binary recommendation experiment satisfies monotone likelihood ratio order in θ, and*

$$\bar{\mu}_\iota(m, 1) \ge \bar{\mu}_\iota(m, 0)$$

for every message m.

Proof. Write $p(\theta) := p_\iota(\theta) = 1 - \iota G_\varepsilon(s_0 - \theta)$. If g_ε is the density of G_ε, then

$$p'(\theta) = \iota g_\varepsilon(s_0 - \theta) \ge 0,$$

with strict inequality on interior points where $g_\varepsilon > 0$. For a binary signal, MLR is equivalent to monotonicity of the odds ratio:

$$\frac{\Pr(y = 1 \mid \theta)}{\Pr(y = 0 \mid \theta)} = \frac{p(\theta)}{1 - p(\theta)},$$

which is increasing in θ because $p(\theta)$ is increasing and $x \mapsto x/(1 - x)$ is increasing on $(0, 1)$.

Fix message m and let $\tilde{f}_m$ be the posterior density after observing m alone. Bayes' rule gives

$$\bar{\mu}_\iota(m, 1) - \mathrm{E}_{\tilde{f}_m}[\theta] = \frac{\mathrm{Cov}_{\tilde{f}_m}\big(\theta, p(\theta)\big)}{\mathrm{E}_{\tilde{f}_m}[p(\theta)]} \ge 0,$$

since $p(\theta)$ is increasing in θ. Also

$$\bar{\mu}_\iota(m,0) - \mathrm{E}_{\tilde{f}_m}[\theta] = -\frac{\mathrm{Cov}_{\tilde{f}_m}\big(\theta, p(\theta)\big)}{1-\mathrm{E}_{\tilde{f}_m}[p(\theta)]} \le 0.$$

Hence $\bar{\mu}_\iota(m,1) \ge \bar{\mu}_\iota(m,0)$. □

Lemma 10 (Higher committee independence is Blackwell more informative)**.** *If $0 \le \iota_1 < \iota_2 \le 1$, then the recommendation experiment at ι_2 Blackwell-dominates the experiment at ι_1.*

Proof. Take $0 \le \iota_1 < \iota_2 \le 1$ and define $\rho := \iota_1/\iota_2 \in [0,1)$. Let Y_2 be the recommendation under ι_2, and define a garbling map K from $\{0,1\}$ to $\{0,1\}$:

$$K(1 \mid 1) = 1, \quad K(0 \mid 1) = 0, \quad K(0 \mid 0) = \rho, \quad K(1 \mid 0) = 1-\rho.$$

So the garbled signal Y_1 equals Y_2 when $Y_2 = 1$, and flips $0 \to 1$ with probability $1-\rho$ when $Y_2 = 0$. Then

$$\Pr(Y_1 = 1 \mid \theta) = \Pr(Y_2 = 1 \mid \theta) + (1-\rho)\Pr(Y_2 = 0 \mid \theta) = (1-\rho) + \rho\, p_{\iota_2}(\theta) = 1 - \iota_1 G_\varepsilon(s_0 - \theta) = p_{\iota_1}(\theta).$$

Thus the ι_1 experiment is a garbling of the ι_2 experiment. Therefore ι_2 Blackwell-dominates ι_1. □

Proposition 21 (Required persuasion message by recommendation state)**.** *Fix a voting bar u and suppose $m \mapsto \mu_\iota(m,y)$ is strictly increasing for each y. Define*

$$m_\iota^{\mathrm{req}}(u,y) := \inf\{m : \mu_\iota(m,y) \ge u\}.$$

Then

$$m_\iota^{\mathrm{req}}(u,1) \le m_\iota^{\mathrm{req}}(u,0).$$

Hence favorable committee recommendations weakly lower the minimum persuasive message needed to clear a fixed voting bar.

Proof. Fix (u,ι). By Lemma 9,

$$\bar{\mu}_\iota(m,1) \ge \bar{\mu}_\iota(m,0) \quad \forall m.$$

Since $\mu_\iota(m,y)$ is an affine map of $\bar{\mu}_\iota(m,y)$ with the same weight $\tau(m,y) \in [0,1]$, the order is preserved:

$$\mu_\iota(m,1) \ge \mu_\iota(m,0) \quad \forall m.$$

Let $m_0 := m_\iota^{\mathrm{req}}(u,0)$. By definition, $\mu_\iota(m_0,0) \ge u$. Hence $\mu_\iota(m_0,1) \ge u$, so m_0 belongs to the feasible set defining $m_\iota^{\mathrm{req}}(u,1)$. Therefore

$$m_\iota^{\mathrm{req}}(u,1) \le m_0 = m_\iota^{\mathrm{req}}(u,0).$$

□

Proposition 22 (Signal channel and enforcement exposure at fixed message)**.** *Fix (m, r, σ) and define*

$$\mathcal{E}(\iota; m) := \int_{\underline{\theta}_0}^{\overline{\theta}_0} L\big((m-\theta)_+, r, \sigma\big)\, p_\iota(\theta)\, f(\theta)\, d\theta.$$

Then $\mathcal{E}(\iota; m)$ is weakly decreasing in ι.

Proof. Using

$$p_\iota(\theta) = 1 - \iota G_\varepsilon(s_0 - \theta),$$

we have

$$\frac{\partial p_\iota(\theta)}{\partial \iota} = -G_\varepsilon(s_0 - \theta) \le 0.$$

Differentiate under the integral sign:

$$\frac{\partial \mathcal{E}(\iota; m)}{\partial \iota} = \int_{\underline{\theta}_0}^{\overline{\theta}_0} L\big((m-\theta)_+, r, \sigma\big)\, \frac{\partial p_\iota(\theta)}{\partial \iota}\, f(\theta)\, d\theta$$

$$= -\int_{\underline{\theta}_0}^{\overline{\theta}_0} L\big((m-\theta)_+, r, \sigma\big)\, G_\varepsilon(s_0 - \theta)\, f(\theta)\, d\theta \le 0,$$

because $L(\cdot, r, \sigma) \ge 0$, $G_\varepsilon \ge 0$, and $f \ge 0$. Therefore $\mathcal{E}(\iota; m)$ is weakly decreasing in ι. □

E Appendix D: Welfare extension proofs and closed-form computations

E.1 Proof of Proposition 4

Proof. By definition,

$$W^{\text{firm}}(\pi) = \int_{\theta^*(\pi)}^{\overline{\theta}_0} \theta f(\theta)\, d\theta - \omega\, \text{EC}(\pi), \qquad W^{\text{FB}} = \int_0^{\overline{\theta}_0} \theta f(\theta)\, d\theta.$$

Hence

$$W^{\text{FB}} - W^{\text{firm}}(\pi) = \int_0^{\overline{\theta}_0} \theta f(\theta)\, d\theta - \int_{\theta^*(\pi)}^{\overline{\theta}_0} \theta f(\theta)\, d\theta + \omega\, \text{EC}(\pi).$$

Case 1: $\theta^*(\pi) < 0$. Then

$$\int_{\theta^*(\pi)}^{\overline{\theta}_0} \theta f(\theta)\, d\theta = \int_{\theta^*(\pi)}^{0} \theta f(\theta)\, d\theta + \int_0^{\overline{\theta}_0} \theta f(\theta)\, d\theta,$$

so

$$W^{\mathrm{FB}} - W^{\mathrm{firm}}(\pi) = -\int_{\theta^*(\pi)}^{0} \theta f(\theta)\, d\theta + \omega\, \mathrm{EC}(\pi) = \mathrm{FP}(\pi) + \omega\, \mathrm{EC}(\pi),$$

and $\mathrm{FN}(\pi) = 0$ by definition.

Case 2: $\theta^*(\pi) > 0$. Then

$$\int_{\theta^*(\pi)}^{\overline{\theta}_0} \theta f(\theta)\, d\theta = \int_{0}^{\overline{\theta}_0} \theta f(\theta)\, d\theta - \int_{0}^{\theta^*(\pi)} \theta f(\theta)\, d\theta,$$

so

$$W^{\mathrm{FB}} - W^{\mathrm{firm}}(\pi) = \int_{0}^{\theta^*(\pi)} \theta f(\theta)\, d\theta + \omega\, \mathrm{EC}(\pi) = \mathrm{FN}(\pi) + \omega\, \mathrm{EC}(\pi),$$

and $\mathrm{FP}(\pi) = 0$.

Case 3: $\theta^*(\pi) = 0$. Then $\mathrm{FP}(\pi) = \mathrm{FN}(\pi) = 0$ and

$$W^{\mathrm{FB}} - W^{\mathrm{firm}}(\pi) = \omega\, \mathrm{EC}(\pi).$$

Combining all cases yields

$$W^{\mathrm{firm}}(\pi) = W^{\mathrm{FB}} - \mathrm{FP}(\pi) - \mathrm{FN}(\pi) - \omega\, \mathrm{EC}(\pi).$$

□

E.2 Proof of Proposition 5

Proof. Under Assumptions 4-6,

$$\mathrm{EC}(\pi) = \int_{\theta^*(\pi)}^{u(\pi)} \phi(\theta, \pi)\, d\theta, \quad \phi(\theta, \pi) = L(m^{\mathrm{req}}(u(\pi)) - \theta, r(\pi), \sigma)\, f(\theta).$$

Leibniz' rule gives

$$\frac{d}{d\pi}\mathrm{EC}(\pi) = \phi(u, \pi) u'(\pi) - \phi(\theta^*, \pi)\theta^{*\prime}(\pi) + \int_{\theta^*}^{u} \partial_\pi \phi(\theta, \pi)\, d\theta.$$

Now

$$\partial_\pi \phi(\theta, \pi) = \left[L_g(m^{\mathrm{req}}(u) - \theta, r(\pi), \sigma)\, \frac{d}{d\pi} m^{\mathrm{req}}(u(\pi)) + L_r(m^{\mathrm{req}}(u) - \theta, r(\pi), \sigma)\, r'(\pi) \right] f(\theta),$$

and

$$\frac{d}{d\pi} m^{\mathrm{req}}(u(\pi)) = (m^{\mathrm{req}})'(u) u'(\pi).$$

Substituting yields the first derivative formula in the proposition.

For welfare,

$$W^{\text{firm}}(\pi) = \int_{\theta^*(\pi)}^{\overline{\theta}_0} \theta f(\theta)\, d\theta - \omega\, \mathrm{EC}(\pi),$$

so again by Leibniz:

$$\frac{d}{d\pi} W^{\text{firm}}(\pi) = -\theta^*(\pi) f(\theta^*(\pi)) \theta^{*\prime}(\pi) - \omega \frac{d}{d\pi}\mathrm{EC}(\pi).$$

If $\theta^*(u, r)$ is C^1, then

$$\theta^{*\prime}(\pi) = \theta^*_u(u(\pi), r(\pi)) u'(\pi) + \theta^*_r(u(\pi), r(\pi)) r'(\pi).$$

Substitute this into the expression for $\mathrm{EC}'(\pi)$ and collect coefficients of $u'(\pi)$ and $r'(\pi)$:

$$\frac{d}{d\pi}\mathrm{EC}(\pi) = u'(\pi)\mathcal{D}_u(\pi) + r'(\pi)\mathcal{D}_r(\pi),$$

where

$$\mathcal{D}_u(\pi) := \phi(u, \pi) - \phi(\theta^*, \pi)\theta^*_u(u(\pi), r(\pi)) + \int_{\theta^*}^{u} L_g(m^{\text{req}}(u) - \theta, r(\pi), \sigma)\, (m^{\text{req}})'(u) f(\theta)\, d\theta,$$

$$\mathcal{D}_r(\pi) := -\phi(\theta^*, \pi)\theta^*_r(u(\pi), r(\pi)) + \int_{\theta^*}^{u} L_r(m^{\text{req}}(u) - \theta, r(\pi), \sigma)\, f(\theta)\, d\theta.$$

This is the claimed decomposition. □

E.3 Proof of Proposition 6

Proof. Under identity beliefs, $m^{\text{req}}(u) = u$ and for pooling types $\theta \in [\theta^*, u)$:

$$g = u - \theta, \qquad w := u - \theta^*.$$

At the boundary $\theta^* = u - w$, indifference $\Lambda(\theta^*; u) = 0$ implies

$$\alpha\theta^* + b_0 + b_1\theta^* - \eta(r, \sigma)(u - \theta^*) - \frac{\kappa(r, \sigma)}{2}(u - \theta^*)^2 = 0,$$

hence

$$\frac{\kappa(r, \sigma)}{2} w^2 + \big(\alpha + b_1 + \eta(r, \sigma)\big) w = (\alpha + b_1)u + b_0.$$

For $\kappa(r, \sigma) > 0$, solving the quadratic yields the positive root

$$w(u, r) = \frac{-\big(\alpha + b_1 + \eta(r, \sigma)\big) + \sqrt{\big(\alpha + b_1 + \eta(r, \sigma)\big)^2 + 2\kappa(r, \sigma)\big((\alpha + b_1)u + b_0\big)}}{\kappa(r, \sigma)}.$$

For $\kappa(r, \sigma) = 0$, the equation is linear and gives

$$w(u, r) = \frac{(\alpha + b_1)u + b_0}{\alpha + b_1 + \eta(r, \sigma)}.$$

Therefore $\theta^*(u, r) = u - w(u, r)$.

With uniform density $f(\theta) = 1/\Delta_\theta$ on $[\underline{\theta}_0, \overline{\theta}_0]$, $\Delta_\theta = \overline{\theta}_0 - \underline{\theta}_0$:

$$\mathrm{EC}(u, r) = \frac{1}{\Delta_\theta} \int_{\theta^*}^{u} \left[\eta(r, \sigma)(u - \theta) + \frac{\kappa(r, \sigma)}{2}(u - \theta)^2 \right] d\theta.$$

Set $t = u - \theta$, so t runs from 0 to w:

$$\mathrm{EC}(u, r) = \frac{1}{\Delta_\theta} \left[\eta(r, \sigma) \int_0^w t\, dt + \frac{\kappa(r, \sigma)}{2} \int_0^w t^2\, dt \right] = \eta(r, \sigma) \frac{w^2}{2\Delta_\theta} + \kappa(r, \sigma) \frac{w^3}{6\Delta_\theta}.$$

Firm-value welfare is

$$W^{\mathrm{firm}}(u, r) = \frac{1}{\Delta_\theta} \int_{\theta^*(u,r)}^{\overline{\theta}_0} \theta\, d\theta - \omega\, \mathrm{EC}(u, r) = \frac{\overline{\theta}_0^{\,2} - \theta^*(u, r)^2}{2\Delta_\theta} - \omega\, \mathrm{EC}(u, r).$$

All claimed expressions follow. □

E.4 Deferred derivations from the main text

This subsection collects material moved from the main text to streamline the exposition. Cross-references in Sections 4–5 point here.

E.4.1 Proof of Lemma 1 (ratification cost)

Proof. Step 1: vote cutoff and feasibility. If the controller buys all voters with $v \le v^*$, the purchased fraction is $G(v^*)$. The credited fraction is $(1 - \sigma)G(v^*)$. Meeting threshold r requires

$$(1 - \sigma)\, G(v^*) \;=\; r \qquad \Longrightarrow \qquad v^* \;=\; G^{-1}\!\left(\frac{r}{1 - \sigma}\right).$$

Since $G(\overline{v}) = 1$, feasibility requires $r/(1 - \sigma) \le 1$, i.e. $r \le 1 - \sigma$. If $r > 1 - \sigma$, then $v^* > \overline{v}$ and ratification is infeasible, so $C = \infty$. Below we work on $r \le 1 - \sigma$.

Step 2: total expenditure. The minimum transfer for voter type v is $t(v) = (R + v)_+$. The total expenditure is therefore

$$C \;=\; \int_{\underline{v}}^{v^*} (R + v)_+\, g(v)\, dv \;=\; \int_{\max\{\underline{v}, -R\}}^{v^*} (R + v)\, g(v)\, dv,$$

where the second equality uses $(R+v)_+ = 0$ for $v < -R$. If $v^* \le -R$ (equivalently, $G(v^*) \le G(-R)$, i.e. $r \le (1-\sigma)G(-R) = r_0$), the integration interval is empty and $C = 0$. If $v^* > -R$ (i.e. $r > r_0$), we are on the active branch and $C > 0$.

Step 3: partial derivatives. Throughout, write $d := 1 - \sigma \in (0, 1]$ and work on the active interior $r_0 < r < d$ where $v^* > -R$. Note that $v^* = G^{-1}(r/d)$, so

$$\frac{\partial v^*}{\partial r} = \frac{1}{d\, g(v^*)}, \qquad \frac{\partial v^*}{\partial \sigma} = \frac{r}{d^2\, g(v^*)}. \tag{6}$$

Because $g(v^*) > 0$ on the interior, both partials are well defined and strictly positive. Crucially, v^* does not depend on R.

(a) C_R. Write $\ell(R) := \max\{\underline{v}, -R\}$ for the effective lower limit. Differentiate $C = \int_{\ell(R)}^{v^*}(R+v)\, g(v)\, dv$ with respect to R. By Leibniz' rule, v^* does not depend on R, so there is no upper-boundary contribution. The integrand derivative is $\partial(R+v)/\partial R = 1$. If $-R > \underline{v}$, the lower limit is $\ell = -R$ and the boundary evaluation is $(R + (-R))\, g(-R) \cdot (-1) = 0$. Therefore in all cases:

$$C_R = \int_{\ell(R)}^{v^*} g(v)\, dv = G(v^*) - G(\ell(R)) > 0,$$

since $v^* > \ell(R)$ on the active region.

(b) C_{RR}. Differentiate $C_R = G(v^*) - G(\ell(R))$ with respect to R. Since v^* is independent of R, only the lower-limit term contributes. If $-R > \underline{v}$, then $\ell(R) = -R$ and $\partial\ell/\partial R = -1$, giving $C_{RR} = g(-R) > 0$. If $-R \le \underline{v}$, then $\ell(R) = \underline{v}$ is constant in R, so $C_{RR} = 0$. Hence

$$C_{RR} = g(-R)\, \mathbf{1}\{-R > \underline{v}\} \ge 0.$$

(c) C_r. By Leibniz' rule at the upper limit (the lower limit does not depend on r):

$$C_r = (R+v^*)\, g(v^*)\, \frac{\partial v^*}{\partial r} = (R+v^*)\, g(v^*)\, \frac{1}{d\, g(v^*)} = \frac{R+v^*}{d} > 0,$$

where the density $g(v^*)$ cancels. The strict inequality uses $R + v^* > 0$ on the active region.

(d) C_σ. Similarly, by Leibniz at the upper limit:

$$C_\sigma = (R+v^*)\, g(v^*)\, \frac{\partial v^*}{\partial \sigma} = (R+v^*)\, g(v^*)\, \frac{r}{d^2\, g(v^*)} = \frac{r\,(R+v^*)}{d^2} > 0.$$

(e) C_{Rr}. Differentiate $C_r = (R+v^*)/d$ with respect to R. Since v^* and d are independent of R:

$$C_{Rr} = \frac{1}{d} = \frac{1}{1-\sigma} > 0.$$

(f) $C_{R\sigma}$. Differentiate $C_R = G(v^*) - G(\ell(R))$ with respect to σ. Only v^* depends on σ:

$$C_{R\sigma} \;=\; g(v^*)\,\frac{\partial v^*}{\partial \sigma} \;=\; g(v^*)\,\frac{r}{d^2\,g(v^*)} \;=\; \frac{r}{d^2} \;=\; \frac{r}{(1-\sigma)^2} \;>\; 0.$$

Kink. At $r = r_0(R, \sigma)$, we have $v^* = -R$ and $R + v^* = 0$, so $C_r = 0$ and $C_\sigma = 0$. Right-derivatives in r are zero, and all stated inequalities hold weakly. □

E.4.2 Proof of Proposition 1

Proof. We use the sign conditions from Lemma 1. On the active branch $r > r_0(R, \sigma)$:

$$C_R > 0, \quad C_{RR} \geq 0, \quad C_r > 0, \quad C_\sigma > 0, \quad C_{Rr} > 0, \quad C_{R\sigma} > 0. \tag{7}$$

When $r \leq r_0$, $C = 0$ and one-sided derivatives are weakly nonnegative. Recall that $C_{RR} = g_v(-R)\,\mathbf{1}\{-R > \underline{v}\} \geq 0$, where g_v denotes the resistance density; all other signs in (7) are distribution-free (Remark E.4.4).

Verification of (L1). At $g = 0$: $L(0, r, \sigma) = p(0, r)\,C(R(0), r, \sigma) = 0$ by assumptions (a)–(b) ($p(0, r) = 0$ and $R(0) = 0$).

For $L_g \geq 0$, differentiate the product $L = p\,C$:

$$L_g \;=\; p_g\,C \;+\; p\,C_R\,R_g \;\geq\; 0,$$

since $p_g \geq 0$ (assumption (a)), $C \geq 0$, $p \geq 0$, $C_R \geq 0$ (Lemma 1(a)), and $R_g \geq 0$ (assumption (b)).

For $L_{gg} \geq 0$, differentiate again:

$$L_{gg} \;=\; p_{gg}\,C \;+\; 2\,p_g\,C_R\,R_g \;+\; p\,\left(C_{RR}\,R_g^2 \;+\; C_R\,R_{gg}\right) \;\geq\; 0,$$

using $p_{gg} \geq 0$ (assumption (a)), $C_{RR} \geq 0$ (Lemma 1(b)), and $R_{gg} \geq 0$ (assumption (b)). Each of the four terms is individually nonnegative.

Verification of (L2). For $L_r \geq 0$:

$$L_r \;=\; p_r\,C \;+\; p\,C_r \;\geq\; 0,$$

using $p_r \geq 0$ (assumption (a)) and $C_r \geq 0$ (Lemma 1(c)).

For $L_{gr} \geq 0$:

$$L_{gr} \;=\; p_{gr}\,C \;+\; p_g\,C_r \;+\; p_r\,C_R\,R_g \;+\; p\,C_{Rr}\,R_g \;\geq\; 0,$$

using $p_{gr} \geq 0$ (assumption (a)) and $C_{Rr} \geq 0$ (Lemma 1(e)). Each term is individually nonnegative.

Verification of (L3). Since p and R do not depend directly on σ:

$$L_\sigma \;=\; p\,C_\sigma \;\geq\; 0,$$

using $C_\sigma \geq 0$ (Lemma 1(d)). And:

$$L_{g\sigma} \;=\; p_g\,C_\sigma \;+\; p\,C_{R\sigma}\,R_g \;\geq\; 0,$$

using $C_{R\sigma} \geq 0$ (Lemma 1(f)).

Thus (L1)–(L3) hold on the feasible set $0 < r \leq 1-\sigma$ under any continuous resistance distribution with positive density. Weak inequalities obtain at the kink $r = r_0$; strict inequalities hold on the active interior whenever at least one corresponding primitive slope is strictly positive. □

E.4.3 Closed-form ratification cost under uniform resistance

Corollary 5 (Closed-form ratification cost under uniform resistance)**.** *If* $v_i \sim \mathrm{Unif}[-1,1]$*, then* $G(v) = (1+v)/2$*,* $g(v) = 1/2$*,* $v^* = 2r/(1-\sigma)-1$*, and the kink threshold is* $r_0(R,\sigma) = (1-\sigma)(1-R)/2$*. On the feasible region* $r \leq 1-\sigma$*,*

$$C(R,r,\sigma) \;=\; \frac{\big(2r-(1-\sigma)(1-R)\big)_+^2}{4(1-\sigma)^2}.$$

On the active interior,

$$C_R = \frac{Z}{2d} > 0, \qquad C_{RR} = \frac{1}{2} > 0,$$
$$C_r = \frac{Z}{d^2} > 0, \qquad C_\sigma = \frac{rZ}{d^3} > 0,$$

and

$$C_{Rr} = \frac{1}{d} > 0, \qquad C_{R\sigma} = \frac{r}{d^2} > 0,$$

where $Z := 2r-(1-\sigma)(1-R) > 0$ *and* $d := 1-\sigma$*.*

Proof. Substitute $G(v) = (1+v)/2$ and $g(v) = 1/2$ into Lemma 1. The cutoff is $v^* = 2r/d - 1$ and $R + v^* = R - 1 + 2r/d = Z/d$. On the active branch:

$$C = \frac{1}{2}\int_{-R}^{v^*} (R+v)\,dv = \frac{1}{4}(R+v^*)^2 = \frac{Z^2}{4d^2}.$$

The derivative expressions follow from Lemma 1 with $R+v^* = Z/d$, or directly from $C = Z^2/(4d^2)$. □

E.4.4 Distribution-free signs of the ratification cost

Remark (Distribution-free signs). The sign conditions $C_r > 0$, $C_\sigma > 0$, $C_{Rr} > 0$, and $C_{R\sigma} > 0$ hold for any continuous resistance distribution G with positive density. The density $g(v^*)$ cancels in every Leibniz evaluation at the upper limit $v^*(r, \sigma)$, and the cancellation is not a coincidence—it reflects the economics of procurement along a supply curve.

When the ratification threshold rises by dr, the controller must buy a thin additional slice of shareholders near the current pivotal voter v^*. If the density is high at v^* (many shareholders clustered near the margin), each one is barely more expensive than the last, but a large number must be purchased to cover the increment. If the density is low (few shareholders near the margin), each one is substantially more expensive, but only a few are needed. In both cases the total incremental cost is the same: $(R+v^*)\,dr/(1-\sigma)$. The number of additional voters, $g(v^*)\cdot dv^*$, and the rate at which the cutoff moves, $\partial v^*/\partial r = 1/\big((1-\sigma)\,g(v^*)\big)$, always offset. What determines the marginal cost is the *price* at the margin—the pivotal voter's reservation value $R + v^*$—not the *density* at the margin.

The exception is $C_{RR} = g(-R)\,\mathbf{1}\{-R > \underline{v}\}$, which asks a different economic question: not "what does the next vote cost at the top of the supply curve?" but "how many shareholders just crossed from the zero-cost region into the active region as R increased?" This is a question about the density at the *kink*, not the density at the *pivotal voter*, so no offsetting inverse-CDF derivative arises. The curvature of the cost in recovery stakes is the one property that depends on the distributional shape of shareholder heterogeneity.

E.4.5 Primitive beliefs and Bayesian updating at the pooling message

Remark (Primitive beliefs and Bayesian updating at the pooling message). The voting outcome in this model is governed by the *primitive* belief function μ satisfying (M1), not by the equilibrium Bayesian posterior that a fully strategic receiver would compute. Individual minority shareholders are atomistic: each observes the disclosed message m, forms the belief $\mu(m)$, and votes according to the aggregation rule (V2). They do not reverse-engineer the controller's equilibrium strategy to update further.

This distinction matters for the partial-pooling equilibrium of Proposition 2. Types $[\theta^*(u), u)$ pool at $m_{\text{req}}(u)$, and types $\theta \geq u$ separate truthfully. A receiver who solved the full signalling game and observed $m = m_{\text{req}}(u)$ would compute $\mathrm{E}[\theta \mid \theta \in [\theta^*(u), u)] < u$ and vote to reject. However, in the atomistic-voter framework of Levit et al. (2024), the voting input is $\mu(m)$—the primitive effective posterior from (M1)—and the approval condition reduces to the message threshold $m \geq m_{\text{req}}(u) := \mu^{-1}(u)$. Since the pooling message equals $m_{\text{req}}(u)$ by construction, the threshold is met and the proposal is approved.

When attention is partial ($\tau < 1$ in the notation of Section 5.1.3), the effective posterior $\mu(m) = \tau\,\mathrm{E}[\theta \mid m] + (1-\tau)\mathrm{E}[\theta]$ is pulled toward the prior mean, so for sufficiently low τ the effective belief at the pooling message can exceed u even under full Bayesian updating. The benchmark

$\tau = 1$ is the limiting case where the atomistic-voter interpretation is needed.

E.4.6 Proof of Corollary 1 (sharp LA-DWBR under ratification microfoundation)

Proof. Under the microfoundation, the enforcement cost is $L(g, r, \sigma) = p(g, r) \cdot C\big(R(g, x^*), r, \sigma\big)$. Differentiating with respect to r:

$$L_r \;=\; p_r \cdot C \;+\; p \cdot C_r.$$

By Lemma 1(c), on the active branch $r > r_0(R, \sigma)$:

$$C_r(R, r, \sigma) \;=\; \frac{R + v^*(r, \sigma)}{1 - \sigma} \;>\; 0,$$

where $v^*(r, \sigma) = G^{-1}\big(r/(1 - \sigma)\big)$. Substituting into Proposition 3(iii), with $r'(\pi) = 1$, yields (4). At the kink $r = r_0$ we have $R + v^* = 0$, so $C_r = 0$ and the inequality holds weakly. □

F Online Appendix C: Robustness extensions

This appendix collects robustness extensions referenced in the main text. Each subsection formalizes an alternative specification or generalization; none is required for the baseline results.

F.1 Endogenous cleansing scrutiny: $\sigma = \sigma(g)$

The baseline analysis treats court-credit risk σ as exogenous. As a tractable strengthening extension, let cleansing scrutiny depend on the disclosure gap. On the up-talk wedge, where $g \geq 0$, define

$$\sigma \;:\; [0, \infty) \to [0, 1), \qquad \sigma'(g) \geq 0, \qquad \sigma(0) = \sigma_0,$$

and in a two-sided version replace g by $|g|$. The legal interpretation is direct: as the disclosure gap grows, courts are more likely to find that the vote was not fully informed or was effectively coerced, so they are less likely to credit ratification as cleansing.

Define the induced reduced-form enforcement cost

$$\widetilde{L}(g, r) \;:=\; L\big(g, r, \sigma(g)\big).$$

By the chain rule,

$$\frac{d}{dg}\widetilde{L}(g, r) \;=\; L_g\big(g, r, \sigma(g)\big) \;+\; L_\sigma\big(g, r, \sigma(g)\big)\, \sigma'(g).$$

Under Assumption 2, $L_\sigma \geq 0$ and $L_{g\sigma} \geq 0$, so any $\sigma'(g) \geq 0$ steepens the marginal enforcement slope in g, especially for larger gaps. Economically, deception now raises expected enforcement through two channels: direct exposure and a lower probability that courts credit vote-based cleansing. This extension therefore amplifies the litigation channel and makes the LA-DWBR reversal easier to obtain. Relative to a one-stage threshold-bunching benchmark, pooling is less likely to survive once cleansing becomes endogenously harder for larger misrepresentations.

F.2 Continuous-extraction formulation

When extraction is endogenous, the controller chooses $x \in [0, B(\theta)]$ and faces an expected enforcement cost $K(g, x, r, \sigma) = p(g, r) \cdot x + \frac{\varrho}{2} x^2$ with $\varrho > 0$, where $p(g, r)$ is the probability of suit. The optimal extraction is

$$x^*(g, r, \sigma) \;=\; \max\Big\{0,\; \min\Big\{B(\theta),\; \tfrac{1-p(g,r)}{\varrho}\Big\}\Big\},$$

which is continuous in (g, r, σ). Define the reduced-form net cost as $L(g, r, \sigma) := K\big(g, x^*(g, r, \sigma), r, \sigma\big) - x^*(g, r, \sigma)$, so that the controller's Stage-1 objective becomes $\mathbf{1}\{\text{approve}\} \cdot \big(\alpha\theta + (1-\alpha)x^*\big) - K\big(g, x^*, r, \sigma\big)$, which is isomorphic to the fixed-extraction formulation with $L(g, r, \sigma)$ as the cost term. When extraction is fixed at $x = B(\theta)$, the continuous formulation reduces to the lump-sum specification in (C1).

F.3 Alternative settlement-bargaining microfoundation

This subsection provides a reduced-form settlement foundation for $L(g, r, \sigma)$ in the spirit of Bebchuk (1984). Let $p(g)$ be the probability that minority plaintiffs file suit, with $p'(|g|) \geq 0$. Conditional on filing, the controller makes a take-it-or-leave-it settlement offer s (or equivalently chooses settlement intensity), and if settlement is rejected or if ratification is not credited, the case proceeds with expected trial loss $D(g)$ plus litigation cost k, where $D'(|g|) \geq 0$ and $D''(|g|) \geq 0$.

Let $\phi(r, \sigma) \in [0, 1]$ denote the probability that a settlement/ratification package is court-creditable and effectively extinguishes claims, with $\phi_r \leq 0$ and $\phi_\sigma \leq 0$ because stricter vote thresholds and stricter cleansing scrutiny make release harder. A compact expected-cost representation is

$$L(g, r, \sigma) \;=\; p(g)\Big[s(g, r, \sigma) \;+\; \big(1 - \phi(r, \sigma)\big) D(g) \;+\; k\Big].$$

Under standard monotonicity conditions ($p'(|g|) \geq 0$, $s_g \geq 0$, $D'(|g|) \geq 0$, and $\phi_r, \phi_\sigma \leq 0$), L is increasing in $|g|$, increasing in r, and increasing in σ. If p and D are convex in $|g|$ and settlement terms worsen with case strength, L is convex in g. Supermodularity ($L_{gr} \geq 0$, $L_{g\sigma} \geq 0$) follows because larger gaps both raise trial stakes and reduce the practical ability to buy a release through ratification.

F.4 S6 tail-growth: bounded-gap caveat

The signaling cost depends on the gap $\Delta := m - \theta$ via a convex, wedge-symmetric map $C(\Delta)$ with $C(0) = 0$ and C strictly increasing in $|\Delta|$.

(a) *Unbounded-gap case.* If for every θ the feasible message set allows a sequence $\{m_n\}$ with $|m_n - \theta| \to \infty$, then

$$C(\Delta) \xrightarrow[|\Delta|\to\infty]{} \infty,$$

so sufficiently large lies are eventually prohibitive.

(b) *Bounded-gap case (when tail-growth may not bite).* If the feasible set and posteriors imply a *gap envelope*

$$\bar{\Delta}(u) \;=\; m^{\text{req}}(u) - u \quad \text{is bounded on } u \in (\underline{\theta}_0, \overline{\theta}_0),$$

then $|\Delta|$ cannot grow arbitrarily and the "tail-growth" force in (a) need not trigger separation. In this case, separation requires a finite-bar dominance test,

$$\Xi_h(u) \;=\; \alpha u + B(u) - C\big(\bar{\Delta}(u)\big) \;\le\; 0 \quad \text{for some interior } u,$$

rather than appeal to $|\Delta| \to \infty$.

G Bibliography

Rabah Amir. Supermodularity and complementarity in economics: An elementary survey. *Southern Economic Journal*, 71(3):636–660, 2005. 10

Christian At, Sylvain Béal, and Pierre-Henri Morand. Freezeout, compensation rules, and voting equilibria. *International Review of Law and Economics*, 41:91–102, 2015. 4

Vladimir Atanasov, Bernard S. Black, and Conrad S. Ciccotello. Law and tunneling. *Journal of Corporation Law*, 37(2):1–49, 2011. 35

Laurent Bach and Daniel Metzger. How close are close shareholder votes? *Review of Financial Studies*, 32(8):3183–3214, 2019. 2

Lucian A. Bebchuk and Assaf Hamdani. Independent directors and controlling shareholders. *University of Pennsylvania Law Review*, 165(6):1271–1315, 2017. 41

Lucian A. Bebchuk and Kobi Kastiel. The Untenable Case for Perpetual Dual-Class Stock. *Virginia Law Review*, 103:585, 2017. URL https://heinonline.org/HOL/Page?handle=hein.journals/valr103&id=609&div=&collection=. 2

Lucian Arye Bebchuk. Litigation and settlement under imperfect information. *RAND Journal of Economics*, 15(3):404–415, 1984. 5, 13, 14, 84

Lucian Arye Bebchuk and Marcel Kahan. Adverse selection and gains to controllers in corporate freezeouts. In Randall K. Morck, editor, *Concentrated Corporate Ownership*, pages 247–259. University of Chicago Press, 2000. 4

Marco Becht, Andrea Polo, and Stefano Rossi. Does mandatory shareholder voting prevent bad acquisitions? *Review of Financial Studies*, 29(11):3035–3067, 2016. 2, 36

Dipesh Bhattarai, Brian Blank, Tingting Liu, Kathryn Schumann-Foster, and Tracie Woidtke. Proxy exempt solicitation campaigns. Working paper (SSRN), October 2022, 2022. 59

Marcello Bianchi, Luca Enriques, and Mateja Malič. Enforcing rules on related party transactions in Italy: One securities regulator's challenge. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*. Cambridge University Press, 2019. 37

Andreas Blume, Oliver J Board, and Kohei Kawamura. Noisy talk. *Theoretical Economics*, 2(4): 395–440, 2007. 67

Patrick Bolton, Marco Becht, and Ailsa Röell. Corporate governance and control. In George M. Constantinides, Milton Harris, and René M. Stulz, editors, *Handbook of the Economics of Finance*, volume 1, pages 1–109. Elsevier, 2003. 5

Audra L. Boone, Brian J. Broughman, and Antonio J. Macias. Shareholder approval thresholds in acquisitions: Evidence from tender offers. *Journal of Corporate Finance*, 53:225–245, 2018. 2, 36

Alon Brav, Matthew Cain, and Jonathon Zytnick. Retail shareholder participation in the proxy process: Monitoring, engagement, and voting. *Journal of Financial Economics*, 144(2):492–522, May 2022. ISSN 0304-405X. doi: 10.1016/j.jfineco.2021.07.013. URL https://www.sciencedirect.com/science/article/pii/S0304405X2100341X. 16

David Burgstahler and Ilia Dichev. Earnings management to avoid earnings decreases and losses. *Journal of Accounting and Economics*, 24(1):99–126, 1997. doi: 10.1016/S0165-4101(97)00017-7. 2, 59

Mike Burkart, Samuel Lee, and Paul Voss. The evolution of the market for corporate control. *European Corporate Governance Institute–Finance Working Paper*, (956):24–05, 2024. 21

Zhihong Chen, Bin Ke, and Zhifeng Yang. Minority shareholders' control rights and the quality of corporate decisions in weak investor protection countries: A natural experiment from China. *The Accounting Review*, 88(4):1211–1238, 2013. 36

Stijn Claessens, Simeon Djankov, and Larry H.P. Lang. The separation of ownership and control in east asian corporations. *Journal of Financial Economics*, 58(1-2):81–112, 2000. 2

Pierre-Henri Conac, Luca Enriques, and Martin Gelter. Constraining dominant shareholders' self-dealing: The legal framework in France, Germany, and Italy. *European Company and Financial Law Review*, 4(4):491–528, 2007. 36, 37

James D. Cox, Tomas Mondino, and Randall S. Thomas. Understanding the (ir)relevance of shareholder votes on M&A deals. *Duke Law Journal*, 69:503–571, 2019. 4

Ettore Croci, Eric Nowak, and Olaf Ehrhardt. The corporate governance endgame: Minority squeeze-out regulation and post-deal litigation in Germany. *Managerial Finance*, 43(1):95–123, 2017. 36

Ya Dai, Liang Guo, and Palani-Rajan Kadapakkam. Protecting the weak: Efficacy of mandated auctions in minority buyouts. *International Review of Financial Analysis*, 61:9–19, 2019. 35

Elif Dalkır, Mehmet S. Dalkır, and Doron Levit. Freeze-out mergers. *Review of Financial Studies*, 32(8):3266–3297, 2019. 4

Jens Dammann. Related party transactions and intragroup transactions. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*, pages 218–250. Cambridge University Press, 2019. 5

Paul L. Davies. Related party transactions: UK model. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*. Cambridge University Press, 2019. 36

Paul L. Davies. Majority and minority shareholders. In *Introduction to Company Law*, Clarendon Law Series. Oxford University Press, 3rd edition, 2020. 36

Francois Degeorge, Jayendu Patel, and Richard Zeckhauser. Earnings management to exceed thresholds. *The Journal of Business*, 72(1):1–33, 1999. doi: 10.1086/209601. 59

Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. The law and economics of self-dealing. *Journal of financial economics*, 88(3):430–465, 2008. 5, 35

Alexander Dyck and Luigi Zingales. Private benefits of control: An international comparison. *Journal of Finance*, 59(2):537–600, 2004. 5, 35

Luca Enriques and Tobias H. Tröger, editors. *The Law and Finance of Related Party Transactions*. International Corporate Law and Financial Market Regulation. Cambridge University Press, 2019. 5

Luca Enriques and Paolo F. Volpin. Corporate governance reforms in continental Europe. *Journal of Economic Perspectives*, 21(1):117–140, 2007. 37

Michael Ewens, Kairong Xiao, and Ting Xu. Regulatory costs of being public: Evidence from bunching estimation. *Journal of Financial Economics*, 153:103775, 2024. 2

Itai Fiegenbaum. The geography of MFW-land. *Delaware Journal of Corporate Law*, 41:693–754, 2016. 4

Itai Fiegenbaum. The controlling shareholder enforcement gap. *American Business Law Journal*, 56(3):583–644, 2019. 4, 5

Jesse M. Fried, Ehud Kamar, and Yishay Yafeh. The effect of minority veto rights on controller pay tunneling. *Journal of Financial Economics*, 138(3):777–788, 2020. 2, 35

Franklin A. Gevurtz. The shareholder approval conundrum. *Hastings Law Journal*, 70:1105–1162, 2019. 4

Ronald J. Gilson. Controlling shareholders and corporate governance: Complicating the comparative taxonomy. *Harvard Law Review*, 119:1641–1679, 2006a. 5

Ronald J. Gilson. Controlling shareholders and corporate governance: Complicating the comparative taxonomy. *Harvard Law Review*, 119(6):1641–1679, 2006b. 5

Ronald J Gilson and Jeffrey N Gordon. Controlling controlling shareholders. *U. Pa. L. Rev.*, 152: 785, 2003. 5, 41

Ronald J. Gilson and Alan Schwartz. Constraints on private benefits of control: Ex Ante control mechanisms versus Ex Post transaction review. *Journal of Institutional and Theoretical Economics*, 169(1):160–183, 2013. 5

Zohar Goshen. The efficiency of controlling corporate self-dealing: Theory meets reality. *Calif. L. Rev.*, 91:393, 2003. 5

Zohar Goshen and Assaf Hamdani. Corporate control and idiosyncratic vision. *Yale Law Journal*, 125(3):560–617, 2016. 36

Zohar Goshen and Assaf Hamdani. Corporate control and the regulation of controlling shareholders. *The Law and Finance of Related Party Transactions, Luca Enriques & Tobias H. Tröger (Eds.), Cambridge University Press*, 2019. 5

Alperen Afşin Gözlügöl. Blinded by 'Fairness': Why we need (strong) procedural safeguards in screening self-dealing and obtaining a fair price is not the answer. *European Business Organization Law Review*, 23(3):633–669, 2022. 36

Sanford J Grossman and Motty Perry. Perfect sequential equilibrium. *Journal of economic theory*, 39(1):97–119, 1986. 16

María Gutiérrez Urtiaga and María Isabel Sáez Lacave. A contractual approach to disciplining self-dealing by controlling shareholders. *Journal of Law, Finance, and Accounting*, 2(1):173–204, 2017. 5

Assaf Hamdani and Yishay Yafeh. Institutional investors as minority shareholders. *Review of Finance*, 17(2):691–725, 2013. 12, 17

Claire A. Hill and Brett H. McDonnell. Sanitizing interested transactions. *Indiana Law Journal*, 86:1169–1210, 2011. 4

Clifford G. Holderness. Equity issuances and agency costs: The telling story of shareholder approval around the world. *Journal of Financial Economics*, 129(2):415–439, 2018. 2

Henry T. C. Hu and Bernard S. Black. The new vote buying: Empty voting and hidden (morphable) ownership. *Southern California Law Review*, 79(4):811–908, 2006. 12

Ming Jian and T. J. Wong. Propping through related party transactions. *Review of Accounting Studies*, 15(1):70–105, 2010. 36

Navin Kartik. Strategic communication with lying costs. *The Review of Economic Studies*, 76(4): 1359–1395, 2009. 4, 7, 9

Kenneth Khoo and Roberto Tallarita. The price of delaware corporate law reform. *Available at SSRN 5318203*, 2025. 37, 40, 45

Kon Sik Kim. Related party transactions in East Asia. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*. Cambridge University Press, 2019. 35, 36

Henrik Jacobsen Kleven. Bunching. *Annual Review of Economics*, 8(1):435–464, 2016. doi: 10.1146/annurev-economics-080315-015234. 2, 59

Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. Corporate ownership around the world. *Journal of Finance*, 54(2):471–517, 1999. 2

Beni Lauterbach and Yevgeny Mugerman. The effect of institutional investors' voice on the terms and outcome of freeze-out tender offers. *Quarterly Journal of Finance*, 10(1):2050002, 2020. 16

Beni Lauterbach, Evgeny Lyandres, Yevgeny Mugerman, and Barak Yarkoni. The choice between various freeze-out procedures and its consequences. *Journal of Law, Finance, and Accounting*, 6(2):315–351, 2021. 5

Lewis H. Lazarus and Brett M. McCartney. Standards of review in conflict transactions on motions to dismiss: Lessons learned in the past decade. *Delaware Journal of Corporate Law*, 36:967–1012, 2011. 4

David K Levine and Thomas R Palfrey. The paradox of voter participation? a laboratory study. *American political science Review*, 101(1):143–158, 2007. 67

Doron Levit, Nadya Malenko, and Ernst Maug. Trading and shareholder democracy. *The Journal of Finance*, 79(1):257–304, 2024. 17, 19, 82

Elliott R. Levy. Freeze-out transactions the pure way: Reconciling judicial asymmetry between tender offers and negotiated mergers. *Fordham Journal of Corporate & Financial Law*, 11:173–216, 2005. 4

Nan Li. Do majority-of-minority shareholder voting rights reduce expropriation? Evidence from related party transactions. *Journal of Accounting Research*, 59(4):1385–1423, 2021. 2, 35

Amir N. Licht. Farewell to fairness: Towards retiring Delaware's entire fairness review. *Delaware Journal of Corporate Law*, 44:1–64, 2020. ECGI Law Working Paper No. 439/2019. 4

Ernest Lim. *A Case for Shareholders' Fiduciary Duties in Common Law Asia*. International Corporate Law and Financial Market Regulation. Cambridge University Press, 2019. 35

Stefano Lovo, Myron B. Slovin, and Marie E. Sushka. A contracting model of entire fairness. *Journal of Corporate Finance*, 71:102099, 2021. 5

Filip Matějka and Guido Tabellini. Electoral competition with rationally inattentive voters. *Journal of the European Economic Association*, 19(3):1899–1935, 2021. 67

Adam Meirowitz and Shaoting Pi. Voting and trading: The shareholder's dilemma. *Journal of Financial Economics*, 146(3):1073–1096, 2022. 17

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, pages 157–180, 1994. 60

Martin J Osborne, Jeffrey S Rosenthal, and Matthew A Turner. Meetings with costly participation. *American Economic Review*, 90(4):927–943, 2000. 67

Wenjing Ouyang and Pengcheng Zhu. An international study of shareholder protection in freeze-out M&A transactions. *International Review of Financial Analysis*, 45:157–171, 2016. 35

Alessio M. Pacces. Procedural and substantive review of related party transactions. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*, pages 181–217. Cambridge University Press, 2019. 5

Thomas R Palfrey and Howard Rosenthal. A strategic calculus of voting. *Public choice*, 41(1):7–53, 1983. 67

Adam C. Pritchard. Tender offers by controlling shareholders: The specter of coercion and fair price. *Berkeley Business Law Journal*, 1:83–138, 2004. 4

Fernán Restrepo. Judicial deference, procedural protections, and deal outcomes in freezeout transactions: Evidence from the effect of mfw. *Journal of Law, Finance, and Accounting*, 6:353–394, 2021. 2, 33, 36

Fernán Restrepo and Guhan Subramanian. The effect of delaware doctrine on freezeout structure and outcomes: Evidence on the unified approach. *Harvard Business Law Review*, 5(2):205–236, 2015. 4

Emmanuel Saez. Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2(3):180–212, 2010. doi: 10.1257/pol.2.3.180. 2, 59

Andrei Shleifer and Robert W Vishny. Large shareholders and corporate control. *Journal of political economy*, 94(3, Part 1):461–488, 1986. 53

Andrei Shleifer and Robert W. Vishny. A survey of corporate governance. *Journal of Finance*, 52 (2):737–783, 1997. 35

Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3): 665–690, 2003. 67

Guhan Subramanian. Post-Siliconix freeze-outs: Theory and evidence. *Journal of Legal Studies*, 36(1):1–26, 2007. 4

Tobias H. Tröger. Germany's reluctance to regulate related party transactions. In Luca Enriques and Tobias H. Tröger, editors, *The Law and Finance of Related Party Transactions*. Cambridge University Press, 2019. 36

James Si Zeng. Rules versus standards in shareholder protection under Chinese law: A new taxonomy and empirical analysis. *American Journal of Comparative Law*, 2025. Advance article. 36