

Article

Cohesion as Concentration: Exclusion-Driven Fragility in Financial Organizations

Foong Soon Cheong 

Atkinson Graduate School of Management, Willamette University, Salem, OR 97301, USA;
fscheong@willamette.edu

Abstract

Financial crises repeatedly reveal organizations that appear internally aligned while failing to recognize accumulating tail risks. This paper argues that cohesion is observationally ambiguous. It can arise from information integration, in which heterogeneous inputs are debated and synthesized, or from exclusion, in which variance is removed through conformity pressure, gatekeeping, and intolerance of dissent. This distinction is formalized using a signal aggregation model in which an organization maintains an anchor belief and achieves agreement through two exclusion channels: report shrinkage toward the anchor and a tolerance rule that discards reports deviating beyond a threshold. Relative to a full inclusion benchmark, exclusion based cohesion jointly produces state contingent bias that is small in normal regimes but grows sharply under displacement, illusory precision in which observed disagreement falls as tail regime estimation error rises, effective concentration of decision inputs below the nominal participant count, and, when the anchor updates from filtered aggregates, dynamic lock in with delayed regime recognition and abrupt correction. External inputs that bypass internal filtering shorten recognition delays. The model yields testable governance diagnostics linking latent fragility to observable patterns in recorded dissent, anonymous to formal voting gaps, scenario set diversity, pipeline and method concentration, and anchor lag. The central implication is that governance systems should treat low internal conflict and unanimity as potentially diagnostic of variance depletion and should monitor whether heterogeneity is integrated or excluded before stress reveals the difference.

Keywords: exclusion; information aggregation; organizational cohesion; risk governance; tail risk; regime shifts



Academic Editor: Thanasis Stengos

Received: 15 February 2026

Revised: 9 March 2026

Accepted: 11 March 2026

Published: 16 March 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Financial crises and post-crisis governance reviews repeatedly reveal a pattern in which organizations appear internally aligned while failing to recognize material risks. Institutions can exhibit orderly decision processes, stable narratives, and consistent internal reporting even when underlying exposures are misunderstood or underestimated. The collapse of Lehman Brothers illustrates this pattern: formal governance structures and internal reporting frameworks coexisted with valuation practices and exposures that were not fully recognized until the firm entered distress, as documented in the examiner's report (Valukas, 2010). A similar dynamic appeared in the 2023 failure of Silicon Valley Bank, where the Federal Reserve's supervisory review identified weak financial risk management, insufficient independent challenge of interest rate risk assumptions, and delayed corrective action despite existing oversight mechanisms (Federal Reserve Board of Governors, 2023).

In both cases the organizations could appear coherent and well coordinated while critical risk signals failed to meaningfully influence decisions.

The presence of internal agreement therefore provides ambiguous information about organizational health. Agreement can emerge through fundamentally different processes. In one case, organizations collect heterogeneous information, encourage debate, and synthesize diverse inputs into a shared estimate. Through this integration process disagreement is converted into a more accurate representation of the underlying state. In another case, agreement is produced through exclusion. Signals that deviate from dominant assumptions are filtered out through mechanisms such as conformity pressure, narrow recruitment pipelines, methodological gatekeeping, or institutional incentives that discourage dissent. Although both processes can produce similar outward indicators of cohesion, their implications for risk differ sharply. Integration-based cohesion can improve accuracy as informational diversity increases, while exclusion-based cohesion can reduce accuracy precisely in the environments where diversity of signals is most valuable, particularly during regime shifts or tail states in which the true underlying conditions have diverged from prevailing assumptions.

This distinction has practical implications for governance and supervision. Boards, executive committees, and regulators often treat low levels of internal conflict and unanimous decision outcomes as evidence of effective oversight. If agreement is generated through exclusion rather than integration, these indicators may instead reflect a structural vulnerability. An organization may appear stable because dissenting signals are suppressed, not because risks are well understood. In such cases governance assessments risk mistaking the appearance of order for the presence of resilience.

Risk management frameworks currently lack a formal way to distinguish these two sources of cohesion. Without a structured framework, decision makers cannot determine whether internal agreement reflects the successful synthesis of diverse information or the removal of variance from the decision process. The absence of this distinction allows organizations to satisfy procedural governance requirements while carrying latent fragilities that become visible primarily under stress. The central claim of this paper is that apparent cohesion can mask fragility when it arises from exclusion-based filtering that removes heterogeneous signals before they enter the decision aggregate.

To formalize this claim, the paper develops a signal aggregation model in which exclusion operates through an organizational anchor.¹ Individual contributors observe signals about a latent risk state, but their reported views are influenced by conformity pressures that shrink signals toward the anchor and by a tolerance rule that discards reports that deviate beyond a specified threshold. The retained inputs are then aggregated with hierarchical weights that represent institutional authority and influence. This structure captures how organizations can generate agreement through variance reduction rather than through the synthesis of independent information. The model produces several implications. First, filtering generates tail-regime bias: when the true state moves away from the anchor, the signals most informative about the new regime are the most likely to be excluded. Second, the organization can experience illusory precision, in which observed disagreement declines while estimation error increases (Figure 1). Third, filtering and hierarchical weighting reduce the effective number of independent inputs influencing decisions, concentrating informational influence even when nominal participation remains large. Fourth, when the anchor itself updates based on filtered internal estimates, the organization can enter a dynamic lock-in in which recognition of regime shifts is delayed and adjustments occur abruptly once accumulated evidence overwhelms the filtering mechanism (Figure 2).




Figure 1. Exclusion-Driven Illusory Precision. As displacement increases ($|\theta - \bar{s}| / \delta$ rises), observed dispersion among retained reports $D(\delta, \lambda)$ falls while tail-state estimation error increases (shown as tail MSE or $\mathcal{A}_T(\delta, \lambda)$), with divergence becoming pronounced once $|\theta - \bar{s}| / \delta \gtrsim 1$.

The contribution of the paper is a tractable parametric framework that characterizes how exclusion affects the statistical properties of organizational risk estimates relative to a full-inclusion benchmark. The model isolates how tolerance thresholds, conformity pressures, hierarchical aggregation, and anchor updating interact to shape estimation bias, effective independence, and recognition delay. By linking these parameters to governance mechanisms such as hiring pipelines, incentive structures, committee procedures, and methodological norms, the framework provides a way to translate organizational design choices into measurable implications for risk recognition.

Existing conformity and herding models typically derive one or two implications from stylized assumptions. Scharfstein and Stein (1990) derive forecast convergence from reputational incentives; Banerjee (1992) derives cascades from sequential observation; Bénabou (2013) derives collective denial from equilibrium incentive structures. This paper introduces a different analytical object: a single parametric exclusion operator consisting of anchor centered shrinkage combined with a distance based retention rule. This operator jointly produces comparative statics for four outcomes: (i) state contingent tail bias, (ii) illusory precision in which observed disagreement falls while estimation error rises, (iii) effective information concentration below nominal participation, and (iv) anchor lock in with recognition delay under dynamic updating. No existing model links all four properties to a common mechanism with interpretable parameters that map directly to observable governance diagnostics.




Figure 2. Anchor Lock-In and Delayed Recognition. After a regime shift at t_0 (θ_t moves from μ_0 to μ_T), the organizational anchor $\bar{s}_t$ adjusts quickly under integration-based cohesion but exhibits long recognition delay τ and abrupt correction under tight exclusion (low δ , high λ); external variance injection (higher η) shortens the delay.

The paper addresses several related questions. It examines the conditions under which internal agreement reflects resilience rather than variance depletion and considers whether the two can be distinguished before a stress event occurs. It analyzes how exclusion-based filtering distorts aggregate estimates relative to a benchmark in which all signals are incorporated and explores how these distortions scale when the underlying regime shifts. It studies how filtering intensity and hierarchical weight concentration jointly determine the effective number of independent inputs in a decision process. It also analyzes the dynamics of anchor updating to identify conditions under which beliefs become self-reinforcing and to characterize how recognition delay varies with tolerance, conformity pressure, updating speed, and the presence of external informational inputs.

The remainder of the paper develops this argument in several steps. It first positions the analysis within existing studies on groupthink, social learning, governance, and crisis dynamics. It then introduces a conceptual distinction between integration-based and exclusion-based cohesion and defines the mechanisms through which exclusion operates in organizations. The formal model is presented in both static and dynamic forms, beginning with a full-inclusion benchmark and then introducing filtering and anchor updating. Subsequent sections map model parameters to institutional governance mechanisms and derive operational diagnostics that can reveal exclusion-based cohesion before stress occurs. The paper concludes by discussing the limits of the framework and outlining directions for empirical testing and future theoretical extensions.

2. Literature and Positioning

This section reviews four bodies of work that motivate the model: groupthink and conformity, information aggregation under incentives, governance and diversity, and herding and crisis dynamics. The review highlights a gap: existing work describes conformity and monoculture, but often does not provide a unified, parametric link between exclusion intensity and specific tail-risk properties such as state-contingent bias, effective independence, and recognition delay.

2.1. Groupthink, Conformity, and Decision Quality

Janis (1972) introduced groupthink, emphasizing self-censorship, unanimity pressure, and illusion of invulnerability. Subsequent research identifies conditions under which groupthink is more likely, including directive leadership and insulation from external input (Esser, 1998). Work on collective intelligence emphasizes that heterogeneity can improve performance under uncertainty when inputs are effectively aggregated (Page, 2007). Bénabou (2013) formalizes willful blindness and collective denial as equilibrium outcomes under incentives that favor reassurance.

This literature motivates the mechanism but often remains descriptive or does not yield comparative statics that connect exclusion intensity to measurable changes in estimation error across regimes. The present paper captures variance removal as an explicit filtering rule and derives how its effects depend on tolerance, conformity, hierarchy, and updating dynamics.

2.2. Information Aggregation with Incentives and Social Learning

Social learning models show how rational agents can herd when private signals are weak relative to observed actions (Banerjee, 1992; Bikhchandani et al., 1992). In organizations, reputational incentives can induce forecast convergence (Scharfstein & Stein, 1990), and can also generate context-dependent deviations depending on evaluation regimes (Ottaviani & Sørensen, 2006). Organizational architecture affects exploration and exploitation, including how hierarchical versus parallel structures alter accuracy and search (Csaszar & Eggers, 2013).

The present paper complements this literature by modeling exclusion as a distance-based retention rule around a shared anchor, together with conformity-driven report shrinkage and hierarchical aggregation weights. This specification is designed to map directly to governance mechanisms and to generate testable predictions about tail-regime bias and recognition delay.

2.3. Governance, Diversity, and Monitoring

A large body of literature links board composition and governance processes to monitoring intensity and risk outcomes. Evidence suggests that diversity can increase monitoring activity in some settings, though its implications for value and risk depend on institutional context and process design (Adams & Ferreira, 2009). Related work examines how board characteristics correlate with risk measures and outcomes, often with heterogeneous results across samples and identification strategies.

A key gap is that compositional diversity (who is present) does not imply informational inclusion (whose signals enter the decision aggregate with weight). A board can be diverse in observable demographics or credentials while still operating under a narrow tolerance rule that systematically limits the effect of dissenting perspectives. This paper models inclusion explicitly and shows how the effective informational contribution of any member depends on exclusion parameters and weight concentration, not merely on nominal presence.

2.4. Herding, Crowded Positioning, and Crisis Dynamics

Market accounts document correlated exposures and belief convergence as precursors to stress. Bikhchandani and Sharma (2001) distinguish spurious herding from intentional herding, noting that both can produce crowded positioning. Empirical work examines how herding varies across volatility regimes and crisis episodes (Economou et al., 2015). Recent policy discussions emphasize that shared models, third-party dependencies, and AI-related concentration can create correlated reactions to shocks (Financial Stability Board, 2024).

This literature describes downstream consequences of convergence but often does not model upstream organizational processes that generate convergence within firms. The present paper provides an internal mechanism: exclusion around a shared anchor can generate concentrated beliefs and delayed recognition even when underlying information sources are diverse.

2.5. Gap and Novelty

Across these studies, a consistent gap remains: limited integration of (i) a parametric exclusion mechanism, (ii) state-contingent tail error relative to a benchmark, (iii) effective independence under hierarchical weighting, and (iv) recognition delay under anchor updating. This paper fills that gap by modeling exclusion as a distance-based filtering rule and deriving comparative statics that translate into governance diagnostics. The novelty is not the observation that exclusion can degrade decisions, but the structured characterization of when, how, and with what observable signatures.

3. Conceptual Framework: Cohesion Pathways and Exclusion Dimensions

3.1. Two Pathways to Cohesion

Integration-based cohesion preserves variance while enabling coordination. Heterogeneous signals are collected, debated, and synthesized. Disagreement is converted into scenario expansion, conditional commitments, and hedged actions. Mechanisms include structured debate, formal red teams, documented minority views, and psychological safety.

Exclusion-based cohesion reduces variance to produce order. Agreement is achieved by shrinking the space of permissible signals, methods, and participants. Inputs that deviate from an organizational anchor are filtered before they reach the aggregate. Exclusion can be explicit (agenda control, gatekeeping) or implicit (career penalties, self-censorship, methodological legitimacy rules).

Both pathways can produce low overt conflict and stable decision narratives. Their divergence becomes salient under regime shifts, when the true state moves away from the anchor and the most informative signals are those most likely to be filtered.

3.2. Exclusion as a Four-Dimensional Operator

Membership exclusion (people) shapes who occupies decision-relevant roles, affecting the distribution of signal precisions and the prevalence of tail specialists.

Message exclusion (voice) shapes which perspectives are expressed, recorded, and advanced, represented as a tolerance parameter around an anchor.

Method exclusion (evidence) shapes which analytical approaches are treated as legitimate inputs, represented via weight concentration on a subset of estimators and contributor types.

Incentive exclusion (costs) shapes anticipatory conformity and self-censorship, represented as report shrinkage toward the anchor.

These dimensions interact. Strong incentives can substitute for explicit filtering, because self-censoring can reduce how often formal thresholds bind. Narrow membership can make message exclusion less visible, because fewer dissenting signals are generated.

3.3. Observable Equivalence and Latent Divergence

Both pathways can generate similar surface indicators. Latent divergence appears under displacement. Under integration, informative variance remains in the aggregate and pulls estimates toward the new state. Under exclusion, the signals most informative about the new regime are most likely to be filtered, so agreement persists around a state that no longer holds. Figure 1 illustrates the resulting divergence between falling observed dispersion and rising tail-state error as displacement increases. Correction, when it occurs, can be abrupt because the anchor shifts only when filtering is overwhelmed by sufficiently strong forces.

4. Formal Model

This section presents a static aggregation environment, introduces an exclusion mechanism, derives core implications, and then endogenizes the anchor through a dynamic updating rule. (see Appendix A for a consolidated notation reference).

4.1. Environment: Regimes, Organizational Loss, and Assumptions

An organization must estimate a latent risk state θ_t for capital allocation, exposure limits, or strategic commitments. Nature draws θ_t from a two-component mixture that captures normal and tail regimes:

$$\theta_t \sim (1-p)\mathcal{N}(\mu_0, \sigma_0^2) + p\mathcal{N}(\mu_T, \sigma_T^2),$$

with $p \in (0, 1)$, $|\mu_T - \mu_0|$ being large relative to σ_0 , and $\sigma_T^2 \geq \sigma_0^2$.

The organization chooses action a_t based on its estimate $\hat{\theta}_t$ and incurs quadratic loss:

$$L(a_t, \theta_t) = (a_t - \theta_t)^2, \quad a_t = \hat{\theta}_t.$$

The mixture distribution provides a tractable representation of displacement regimes. The model focus is on how exclusion intensity, conformity, hierarchy, and updating dynamics jointly shape tail error and recognition delay, not on the tautology that filtering centered away from the truth creates error.

4.2. Agents, Signals, and the Full-Inclusion Benchmark

There are N potential contributors indexed by $i = 1, \dots, N$. Each observes a private signal:

$$s_{i,t} = \theta_t + \varepsilon_{i,t}.$$

Signals are unbiased. To isolate exclusion effects from correlated information, assume signal errors are conditionally independent across i given θ_t :

$$\varepsilon_{i,t} \perp \varepsilon_{j,t} \mid \theta_t \quad \text{for } i \neq j.$$

Signal noise depends on the regime:

$$\varepsilon_{i,t} \mid \text{normal} \sim \mathcal{N}(0, \sigma_{i,0}^2), \quad \varepsilon_{i,t} \mid \text{tail} \sim \mathcal{N}(0, \sigma_{i,T}^2).$$

A subset $\mathcal{S} \subseteq \{1, \dots, N\}$ are tail specialists, meaning $\sigma_{i,T}^2$ is relatively small for $i \in \mathcal{S}$ compared to the population average.

Oracle Full-Inclusion Benchmark

If all N signals are included and the regime $r \in \{0, T\}$ is known, the precision-weighted mean is:

$$\hat{\theta}_{\text{full},t}^{\text{oracle}} = \frac{\sum_{i=1}^N \omega_{i,r} s_{i,t}}{\sum_{i=1}^N \omega_{i,r}}, \quad \omega_{i,r} = \frac{1}{\sigma_{i,r}^2}.$$

This estimator is unbiased for θ_t and has mean squared error:

$$\text{MSE}_{\text{full},r}^{\text{oracle}} = \frac{1}{\sum_{i=1}^N \omega_{i,r}}.$$

This oracle benchmark is useful for isolating the informational cost of exclusion. A feasible benchmark that does not condition on the realized regime can instead use mixture-averaged precisions or a Bayesian hierarchical update. The qualitative comparisons in this paper do not rely on regime observability, but explicit formulas are cleanest under the oracle benchmark.

4.3. Exclusion Mechanism: Conformity, Tolerance, and Hierarchical Aggregation

The organization maintains an anchor $\bar{s}_t$ that represents an institutional prior, leadership narrative, or dominant internal view. In the static analysis, treat $\bar{s}_t$ as given; Section 4.5 endogenizes it.

4.3.1. Conformity (Report Shrinkage)

Each agent reports:

$$r_{i,t} = (1 - \lambda) s_{i,t} + \lambda \bar{s}_t,$$

where $\lambda \in [0, 1)$ captures conformity pressure and self-censorship. When $\lambda = 0$, reports equal signals.

4.3.2. Message Exclusion (Tolerance Rule)

The organization retains reports within a tolerance band $\kappa > 0$ around the anchor:

$$i \text{ is retained} \Leftrightarrow |r_{i,t} - \bar{s}_t| \leq \kappa.$$

Since $r_{i,t} - \bar{s}_t = (1 - \lambda)(s_{i,t} - \bar{s}_t)$, this is equivalent to a condition on signals:

$$|s_{i,t} - \bar{s}_t| \leq \delta, \quad \delta \equiv \frac{\kappa}{1 - \lambda}.$$

Let $\mathcal{R}_t(\delta)$ be the retained set and $n_t(\delta) = |\mathcal{R}_t(\delta)|$.

4.3.3. Filtered Aggregation with Hierarchy

Retained reports are aggregated using decision weights $w_i \geq 0$ that capture hierarchy, agenda control, or institutional authority:

$$\hat{\theta}_{\text{filt},t}(\delta, \lambda) = \frac{\sum_{i \in \mathcal{R}_t(\delta)} w_i r_{i,t}}{\sum_{i \in \mathcal{R}_t(\delta)} w_i}.$$

4.3.4. Observable Dispersion and Effective Independence

Observed dispersion among retained reports is:

$$D_t(\delta, \lambda) = \text{Var}(r_{i,t} | i \in \mathcal{R}_t(\delta)).$$

A standard measure of effective concentration in weights is:

$$N_{\text{eff},t}(\delta) = \frac{\left(\sum_{i \in \mathcal{R}_t(\delta)} w_i\right)^2}{\sum_{i \in \mathcal{R}_t(\delta)} w_i^2}.$$

When weights are equal over retained agents, $N_{\text{eff},t}(\delta) = n_t(\delta)$. Concentrated weights reduce N_{eff} .

4.4. Core Static Implications

4.4.1. Implication 1: Anchor Pull and Tail Bias

Filtering pulls $\hat{\theta}_{\text{filt},t}$ toward $\bar{s}_t$ through two channels.

Shrinkage channel: Even without selection, the expected report conditional on $s_{i,t}$ is a convex combination of $s_{i,t}$ and $\bar{s}_t$, so the aggregate inherits anchor dependence.

Selection channel: The tolerance rule retains signals within $[\bar{s}_t - \delta, \bar{s}_t + \delta]$. When θ_t is near $\bar{s}_t$, retention is approximately symmetric and bias is small. When θ_t is displaced from $\bar{s}_t$, truncation becomes asymmetric relative to the signal mean and shifts the retained mean toward the anchor.

For Gaussian signals with variance σ_i^2 (suppressing regime index for notation), retention probability is:

$$\mathbb{P}(i \text{ retained} \mid \theta_t, \bar{s}_t) = \Phi\left(\frac{\delta - (\theta_t - \bar{s}_t)}{\sigma_i}\right) - \Phi\left(\frac{-\delta - (\theta_t - \bar{s}_t)}{\sigma_i}\right),$$

where Φ is the standard normal CDF.

Tail specialists and displacement condition. If $|\theta_t - \bar{s}_t| > \delta$, then higher precision (smaller σ_i) can reduce retention because signals cluster near θ_t outside the tolerance band. If $|\theta_t - \bar{s}_t| < \delta$, higher precision increases retention because signals cluster inside the band. The exclusion mechanism is therefore most likely to be accuracy-inverting under displacement, which is the regime of interest for tail-risk governance.

Define tail-regime bias:

$$B_T(\delta, \lambda) = \mathbb{E}[\hat{\theta}_{\text{filt},t}(\delta, \lambda) - \theta_t \mid \text{tail}].$$

Bias increases in magnitude when (i) the anchor is displaced from the tail mean, (ii) tolerance is tighter, and (iii) conformity is stronger.

4.4.2. Implication 2: Illusory Precision

Observed dispersion $D_t(\delta, \lambda)$ among retained reports typically decreases when tolerance tightens or conformity increases, because outlying reports are removed and remaining reports are shrunk toward the anchor. In displaced regimes, the mean squared error can increase:

$$\text{MSE}_T(\delta, \lambda) = \mathbb{E}[(\hat{\theta}_{\text{filt},t}(\delta, \lambda) - \theta_t)^2 \mid \text{tail}].$$

This yields a divergence:

$$D_t(\delta, \lambda) \downarrow \quad \text{while} \quad \text{MSE}_T(\delta, \lambda) \uparrow \quad \text{in displaced tail states.}$$

Define an oracle amplification ratio relative to the oracle full-inclusion benchmark:

$$\mathcal{A}_T(\delta, \lambda) = \frac{\text{MSE}_T(\delta, \lambda)}{\text{MSE}_{\text{full},T}^{\text{oracle}}}.$$

The ratio exceeds 1 when the exclusion rule removes informative signals and reduces effective independence. The joint effect of selection and shrinkage is interactive in the sense that selection-induced bias in the retained mean is scaled by $(1 - \lambda)$ in reports while shrinkage adds an additional anchor term, so increasing conformity can worsen bias even when retention increases.

4.4.3. Implication 3: Concentration Through Retention and Hierarchy

Filtering reduces $n_t(\delta)$ and can change composition by excluding some contributor types at higher rates under displacement. With hierarchical weights, this reduces effective independence:

- Retention loss: fewer retained contributors lowers N_{eff} even under equal weights.
- Composition shift: if filtering changes which contributors remain and weights are concentrated among a subset (such as senior roles), then N_{eff} can fall substantially below $n_t(\delta)$.

A large gap between nominal participation and N_{eff} is an observable signature of concentrated influence.

4.4.4. Implication 4: Common Bias and Coordinated Adjustment (Without Mechanical Correlation)

Under the baseline assumption of conditional independence across agents and an anchor $\bar{s}_t$ treated as given at time t , the term $\lambda \bar{s}_t$ is common but constant across agents, so it does not mechanically create covariance across individual reports conditional on $(\theta_t, \bar{s}_t)$. The primary within-period effect of conformity is therefore cross-sectional compression and shared bias toward $\bar{s}_t$, not induced correlation.

The model nevertheless predicts coordinated adjustment over time because (i) decisions are driven by an anchor-centered filtering rule and (ii) the anchor itself evolves. When the anchor shifts, reported inputs and the filtered aggregate shift in the same direction because the mapping from signals to reports and from reports to inclusion is anchored on the same $\bar{s}_t$. This produces synchronized revisions at the organizational level, such as abrupt changes in limits or positioning, even if underlying signals are not contemporaneously correlated.

If one introduces correlated error components (shared models, vendors, data, or training) or a stochastic anchor with common noise, then the same framework also generates within-period comovement and cross-sectional correlation. That extension is discussed in Section 7.

4.4.5. Workhorse Special Case: Explicit Truncated-Normal Formulas

Consider a static period with equal weights, common signal variance σ^2 , and a fixed anchor $\bar{s}$. Let $s \sim \mathcal{N}(\theta, \sigma^2)$ and retain if $|s - \bar{s}| \leq \delta$. Define

$$a = \frac{\bar{s} - \delta - \theta}{\sigma}, \quad b = \frac{\bar{s} + \delta - \theta}{\sigma}.$$

Then the retained mean of s is:

$$\mathbb{E}[s \mid |s - \bar{s}| \leq \delta] = \theta + \sigma \cdot \frac{\phi(a) - \phi(b)}{\Phi(b) - \Phi(a)},$$

where ϕ is the standard normal PDF. The term $\sigma \cdot \frac{\phi(a) - \phi(b)}{\Phi(b) - \Phi(a)}$ is the selection-induced bias in the retained mean of signals, which is near zero when θ is centered in the band and becomes large in magnitude when θ is displaced relative to $\bar{s}$.

Since $r = (1 - \lambda)s + \lambda\bar{s}$, the retained mean of r is:

$$\mathbb{E}[r \mid |s - \bar{s}| \leq \delta] = (1 - \lambda) \mathbb{E}[s \mid |s - \bar{s}| \leq \delta] + \lambda\bar{s}.$$

Therefore the bias in the retained mean report relative to θ decomposes as:

$$\mathbb{E}[r \mid |s - \bar{s}| \leq \delta] - \theta = (1 - \lambda)(\mathbb{E}[s \mid |s - \bar{s}| \leq \delta] - \theta) + \lambda(\bar{s} - \theta).$$

This expression shows an explicit interaction: selection bias depends on δ through the truncated-normal term, while conformity adds a direct anchor gap term $\lambda(\bar{s} - \theta)$. In displaced regimes, tightening tolerance increases the selection bias component and increasing conformity increases the anchor-gap component, so combined tightening can substantially increase total bias even if conformity increases retention.

4.5. Dynamic Model: Anchor Updating, Lock-In, and Recognition Delay

The static analysis takes $\bar{s}_t$ as given. The dynamic extension endogenizes $\bar{s}_t$ and captures self-reinforcement.

4.5.1. Anchor Updating Rule

Let the anchor update as:

$$\bar{s}_{t+1} = \alpha \hat{\theta}_{\text{filt},t}(\delta, \lambda) + (1 - \alpha - \eta) \bar{s}_t + \eta x_t,$$

where $\alpha \in (0, 1)$ weights the internal filtered estimate, $\eta \in [0, 1)$ weights an external input x_t , and $\alpha + \eta \leq 1$. The external input is not subject to the organization's internal exclusion rule.

4.5.2. Fixed Points and Lock-In

Under a sustained regime with constant $\theta_t = \theta^*$, write $\hat{\theta}_{\text{filt},t} = f(\bar{s}_t; \theta^*, \delta, \lambda)$. A fixed point $\bar{s}^*$ satisfies:

$$\bar{s}^* = \alpha f(\bar{s}^*; \theta^*, \delta, \lambda) + (1 - \alpha - \eta) \bar{s}^* + \eta \mathbb{E}[x_t],$$

equivalently,

$$(\alpha + \eta) \bar{s}^* = \alpha f(\bar{s}^*; \theta^*, \delta, \lambda) + \eta \mathbb{E}[x_t].$$

When $\eta = 0$, fixed points satisfy $\bar{s}^* = f(\bar{s}^*; \theta^*, \delta, \lambda)$. Tight tolerance and strong conformity can make $f(\bar{s})$ close to $\bar{s}$ over wide regions, which supports fixed points far from θ^* and thus lock-in.

4.5.3. Local Stability Condition

Define the one-step mapping:

$$g(\bar{s}) = \alpha f(\bar{s}; \theta^*, \delta, \lambda) + (1 - \alpha - \eta) \bar{s} + \eta \mathbb{E}[x_t].$$

A fixed point $\bar{s}^*$ is locally stable if:

$$|g'(\bar{s}^*)| = \left| \alpha \frac{\partial f}{\partial \bar{s}}(\bar{s}^*; \theta^*, \delta, \lambda) + (1 - \alpha - \eta) \right| < 1.$$

4.5.4. Recognition Delay Under a Regime Shift

Suppose a regime shift occurs at t_0 from a state near μ_0 to a state near μ_T . Define a recognition time τ as the first period after t_0 in which $|\bar{s}_t - \mu_T| \leq \epsilon$ for some tolerance $\epsilon > 0$.

Recognition delay increases when early post-shift filtered estimates remain close to the pre-shift anchor because the tolerance rule excludes signals that point toward the new regime and conformity reduces the information content of those that remain. External inputs can reduce delay by injecting variance that bypasses internal filtering. Faster updating α can reduce delay when internal filtered estimates are informative about the new regime but can also accelerate convergence to an incorrect fixed point when the filtered estimate is heavily anchor-biased. The sign and magnitude of the effect of α depends on how informative $f(\bar{s}_t)$ is in the displacement region.

4.5.5. Abrupt Correction

When internal updating is slow or biased under displacement, the anchor can remain near its prior for extended periods, followed by a relatively rapid transition once enough information enters the process (through external input shocks, overwhelming evidence, or the anchor drifting into a region where the tolerance rule admits more regime-relevant signals). This pattern, extended stability followed by rapid revision, is a signature of exclusion-based cohesion under anchor-centered filtering.

4.6. Comparative Statics Summary

- Tightening exclusion (lower δ) reduces observed disagreement D_t and can increase tail-regime bias and MSE under displacement, reduce N_{eff} , and raise recognition delay.
- Increasing conformity (higher λ) reduces the informational content of reports, increases anchor dependence in aggregates, and can amplify displaced-regime error by adding direct anchor-gap bias.
- Greater hierarchy (more concentrated w_i) lowers N_{eff} and magnifies the effect of composition shifts induced by filtering.
- External variance injection (higher η) can reduce lock-in and recognition delay by bypassing internal exclusion.
- Faster anchor updating (higher α) improves responsiveness when internal filtered estimates are informative, but can worsen lock-in when internal filtering produces strongly anchor-biased estimates in the displacement region.

5. Mechanisms: Institutional Channels as Model Parameters

The parameters in the model correspond to identifiable institutional mechanisms inside financial organizations. The model treats exclusion not as an abstract behavioral tendency but as a set of operational channels that determine which signals are generated, which are expressed, and which ultimately influence decisions. These channels affect the distribution of signals available to the organization, the probability that signals are reported, the likelihood that reported signals are retained in deliberation, and the weight that retained inputs receive in aggregation. In this way, governance structures, incentive systems, and organizational architecture translate directly into the model's parameters governing tolerance, conformity, effective independence, and anchor updating.

Membership exclusion shapes the signal environment by determining who occupies decision-relevant roles and therefore what types of information are produced within the organization. Hiring pipelines, promotion criteria, and professional networks influence the diversity of analytical backgrounds, crisis experience, and epistemic approaches among contributors. When recruitment and advancement reward alignment with dominant internal narratives, heterogeneity in signals tends to decline over time. The result is a narrower distribution of perspectives and a smaller population of contributors whose signals are particularly informative in tail states. Pipeline concentration can also increase correlation in signal errors because individuals trained in similar environments often share analytical

frameworks, data sources, and modeling assumptions. These forms of concentration reduce the effective informational diversity available to the organization even before any explicit filtering occurs.

Message exclusion determines which signals are expressed, recorded, and considered in formal deliberation. In the model this mechanism appears as a tolerance parameter around the organizational anchor that governs which reports enter the decision aggregate. Institutional practices such as committee gatekeeping, agenda control, and escalation procedures can implement effective filtering by restricting which perspectives reach higher levels of the hierarchy. Documentation norms also influence whether dissenting views persist as usable inputs for later updating. When dissent is discouraged or not formally recorded, the informational contribution of divergent signals is lost even if individuals initially generate them. Sequential decision structures can further tighten the effective tolerance band as signals move upward through layers of review, so that perspectives acceptable at lower levels may still be filtered out before reaching final decision makers.

Incentive exclusion operates through the costs associated with expressing views that diverge from prevailing organizational expectations. Career incentives, performance evaluations, and reputational considerations can create conformity pressure that leads individuals to adjust their reported assessments toward the dominant internal narrative. In the model this mechanism appears as report shrinkage toward the anchor. Evaluation systems that reward alignment or penalize visible disagreement increase the incentive to self-censor. Short-horizon compensation structures can intensify this effect because many tail-relevant signals concern low-probability events whose benefits from recognition appear outside standard evaluation windows. Informal sanctions, including reputational penalties or diminished promotion prospects, can reinforce conformity pressures and may contribute to the attrition of contributors who repeatedly provide contrarian assessments.

Method exclusion determines which analytical approaches are treated as legitimate inputs in decision processes and therefore influences the concentration of weights assigned to different contributors or models. Organizations often rely on dominant modeling frameworks, shared data infrastructures, or common vendor systems to generate risk estimates. While such standardization can improve coordination and comparability, it can also concentrate effective decision weights on a narrow set of methodologies. When multiple contributors rely on the same modeling architecture or external provider, their outputs may appear diverse while still sharing underlying assumptions and blind spots. This concentration reduces effective independence in the aggregation process and increases the risk that systematic model errors propagate through the organization without challenge. The marginalization of alternative approaches, including qualitative expertise or historically grounded judgment, can further reduce the influence of contributors who might be most informative in environments characterized by structural change.

External variance injection represents channels through which signals that bypass internal exclusion rules enter the anchor updating process. In the model this appears as an external input parameter that introduces information not filtered through the organization's internal tolerance mechanism. Institutional arrangements such as independent risk review functions, supervisory examinations, or external advisory roles can perform this function when they possess protected authority and direct access to decision forums. These mechanisms can increase the diversity of signals influencing the anchor by introducing perspectives that are less constrained by internal incentives or prevailing narratives. At the same time, the effectiveness of external variance injection depends on governance design. If external inputs are filtered through the same internal structures that shape ordinary reporting, their capacity to offset exclusion-based cohesion may be limited. In contrast, structures that allow external signals to enter the updating process without being fully

mediated by internal anchors can increase responsiveness to regime shifts and reduce the likelihood of prolonged recognition delay.

6. Implications for Governance, Design, and Regulation

The model has several implications for governance design, organizational oversight, and regulatory practice. Because exclusion-based cohesion operates by filtering variance before it enters the decision aggregate, effective governance must monitor the mechanisms that determine whether heterogeneous information is incorporated or removed. Observable organizational indicators can therefore serve as indirect measures of the model's key parameters, particularly tolerance for dissent, conformity pressure, and effective informational independence. Patterns in governance records and internal decision processes may reveal whether agreement reflects genuine information integration or variance depletion produced by exclusion.

One observable indicator concerns the rate at which dissenting views are formally recorded in decision processes. When tolerance for deviation from the organizational anchor narrows, fewer dissenting assessments are likely to appear in the formal record. This diagnostic can be operationalized using committee minutes, board documentation, voting records, and internal memoranda that summarize deliberations. Many financial institutions maintain structured minutes indicating whether decisions were unanimous or whether dissenting views were recorded. A simple empirical measure is the share of decisions with documented dissent relative to the total number of decisions within a committee over time. More granular approaches can code minutes or memoranda for explicit disagreement, minority recommendations, or conditional reservations. A persistent decline in recorded dissent that is not accompanied by a comparable reduction in external uncertainty is consistent with tightening effective tolerance around the anchor.

Another indicator concerns the gap between privately expressed views and formal decisions. If anonymous polling or confidential surveys consistently reveal greater dispersion in opinions than the positions ultimately recorded in official votes, this difference can serve as a proxy for conformity pressure. Empirical implementation can rely on anonymous internal survey instruments, confidential risk culture assessments, or structured pre-vote polling conducted within committees. These sources provide distributions of views that are insulated from reputational exposure. The diagnostic compares dispersion in anonymous responses with dispersion in recorded votes or formal statements. A large and persistent anonymous-to-formal gap suggests that contributors may be adjusting their reported views toward the anchor in anticipation of reputational or career costs associated with visible disagreement.

Additional evidence of exclusion-based cohesion can appear in the evolution of internal scenario sets and analytical frameworks. When organizations repeatedly examine a narrow range of scenarios despite changing external conditions, this narrowing may signal anchor-centered filtering and self-reinforcing internal narratives. Internal scenario documentation provides the primary source for this diagnostic. Risk management units typically maintain libraries of stress scenarios, macroeconomic narratives, and contingency analyses used for capital planning and strategic review. Scenario-set diversity can be evaluated by examining the number of distinct scenarios considered in each review cycle, the dispersion of macroeconomic assumptions embedded in those scenarios, and the conceptual breadth of mechanisms represented in scenario narratives. Text analysis of scenario descriptions can also be used to identify convergence toward a narrow set of explanatory frameworks over time.

A related diagnostic concerns the concentration of analytical methods used in risk estimation. When organizations rely heavily on a single modeling architecture, vendor

platform, or analytical methodology, the effective diversity of estimators entering the decision process declines. Internal model inventories, vendor relationships, and model validation documentation can provide empirical proxies for this form of concentration. Heavy reliance on a small number of analytical frameworks increases the likelihood that systematic model errors propagate through the organization without challenge.

Organizational pipelines provide another observable measure of informational concentration. Hiring and promotion pathways influence the distribution of signal environments represented within decision processes. Pipeline concentration indices can be constructed using human resources data on educational institutions, degree fields, prior employers, and professional certifications. Standard concentration measures can be applied to estimate the share of contributors originating from common training environments or professional networks. Public professional biographies and networking platforms may supplement internal records when direct access to personnel data is limited. High concentration in educational or professional backgrounds suggests that contributors may share similar analytical frameworks and therefore produce more correlated signals.

Recognition delay provides another operational indicator derived from the dynamic model. When external signals suggest a change in underlying conditions but internal scenarios, risk limits, or strategic assumptions only adjust slowly, the lag between external evidence and internal response may reflect anchor lock-in. This diagnostic can be implemented by comparing time series of internal adjustments in risk limits, capital buffers, or scenario assumptions with observable changes in market indicators such as interest rate volatility, credit spreads, macroeconomic forecasts, or regulatory stress scenarios. Internal documentation of risk limit revisions, stress test parameter updates, or portfolio allocation changes provides the internal component of this comparison. Statistical estimates of the lag between external signal shifts and internal responses provide a measure of anchor delay. Prolonged delays followed by abrupt revisions are consistent with the lock-in dynamics predicted by the model.

Turnover among contributors who regularly produce dissenting or tail-oriented analyses provides a related signal. If internal documentation identifies authorship of risk memoranda, scenario analyses, or minority reports, personnel records can be used to examine whether individuals who frequently produce contrarian assessments exit the organization at elevated rates. Such attrition gradually reduces the population of tail specialists and lowers the probability that informative signals about emerging regimes survive the filtering process.

These empirical proxies provide a bridge between the model's theoretical parameters and observable institutional data. Governance records, internal documentation, survey instruments, personnel data, and market time series all contain signals that can be used to infer whether cohesion reflects the integration of heterogeneous information or the systematic exclusion of variance before aggregation.

These diagnostics also suggest several governance design responses. Organizations can create channels that deliberately inject variance into decision processes and thereby counteract exclusion dynamics. Independent risk review functions with protected authority, direct reporting lines to senior oversight bodies, and the capacity to introduce alternative scenarios into strategic discussions can increase the effective weight of external inputs in the anchor updating process. Structured opposition procedures can mitigate conformity pressures by separating information generation from reputational exposure. Mechanisms such as anonymous pre-vote polling, formal challenge rounds, and premortem exercises can expand the range of admissible signals and reduce anticipatory self-censorship. Documenting dissent and revisiting prior minority views during retrospective reviews can

reinforce informational inclusion by demonstrating that dissenting assessments remain part of institutional learning rather than being discarded once a decision is reached.

Organizational hiring and analytical design choices also influence the parameters governing exclusion. Recruiting contributors from outside dominant professional pipelines can expand the distribution of signal environments and increase the likelihood that informative variance enters decision processes. Maintaining multiple analytical approaches and independent model validation processes can reduce method concentration and increase effective informational independence. The objective is not to eliminate convergence entirely, since coordinated action often requires shared estimates, but to ensure that convergence results from the synthesis of heterogeneous inputs rather than pre-aggregation filtering that removes variance before evaluation.

These implications extend to board-level oversight. Boards and senior governance bodies can monitor patterns in dissent rates, anonymous-to-formal opinion gaps, scenario diversity, pipeline concentration, analytical concentration, and recognition delay as indicators of whether internal cohesion reflects informational integration or exclusion-based filtering. Tracking these indicators over time provides a mechanism for detecting latent fragility before stress conditions reveal it through realized losses. Incentive structures can also be aligned with this objective by evaluating leadership not only on narrative stability or short-term performance metrics but also on the organization's capacity to surface and evaluate challenging or tail-oriented assessments.

Regulatory practice may also benefit from incorporating process-focused evaluation alongside traditional capital and liquidity metrics. Supervisory assessments that focus exclusively on formal governance structures may overlook exclusion dynamics if those structures exist but operate under narrow tolerance rules. Evaluating how dissenting analyses enter decisions, how scenarios are generated and revised, and how quickly internal assumptions respond to external evidence can provide additional insight into the resilience of institutional decision processes. At the same time, regulatory design must avoid imposing overly standardized analytical frameworks that unintentionally create cross-firm monoculture. Excessive uniformity in supervisory scenarios or modeling expectations can reduce heterogeneity across institutions and thereby increase the probability of correlated recognition delays and synchronized adjustments during stress episodes.

Taken together, the model suggests that governance systems should treat internal heterogeneity as a risk-management asset rather than a procedural inconvenience. The central task of oversight is therefore not simply to ensure orderly decision processes but to monitor the mechanisms that determine whether diverse signals are incorporated into organizational estimates or filtered out before they can influence outcomes. When cohesion arises from integration, disagreement is transformed into improved understanding of uncertainty. When cohesion arises from exclusion, disagreement disappears from the record while the underlying risk remains.

7. Discussion

This section positions the model within existing research while clarifying what it adds and where its limits lie. The framework is most closely related to work on groupthink, social learning, and organizational governance. However, the contribution here is primarily structural. Rather than describing conformity or consensus qualitatively, the model represents exclusion as a parameterized filtering process centered on a shared organizational anchor. Classic discussions of groupthink emphasize phenomena such as self-censorship, pressure toward unanimity, and insulation from external information. Those insights remain relevant, but they are typically descriptive. The present framework instead treats variance

removal explicitly through tolerance and conformity parameters and examines how these mechanisms affect bias, effective independence, and the timing of regime recognition.

The relationship to social learning models is similar. Standard cascade or herding models explain convergence through sequential imitation or reputational incentives. In contrast, the mechanism here is architectural. Signals are filtered according to their distance from an anchor before they enter the aggregation process. This difference matters because it shifts attention from individual imitation to organizational design. The framework also speaks to the governance literature on board composition and diversity. Much of that work evaluates who sits at the table. The model suggests that composition alone is insufficient. Diversity has informational value only if heterogeneous signals are actually retained and incorporated into the decision aggregate rather than filtered or marginalized.

Like most analytical frameworks, the model relies on several simplifying assumptions. One is the assumption that signal errors are conditionally independent across contributors. In practice, organizational information sources are rarely independent. Shared educational backgrounds, common vendors, and similar analytical frameworks can introduce correlation in errors. Incorporating correlated signals would not alter the core mechanism but would likely strengthen its effects by reducing the effective independence of retained inputs. Another simplification is the use of a binary tolerance rule that either admits or excludes reports. Organizations rarely operate with such a sharp threshold. More commonly they attenuate dissenting views, delay escalation, or give them reduced weight. Modeling these graded forms of filtering would smooth the comparative statics but would leave the underlying logic intact.

The model also treats tolerance and conformity parameters as exogenous. In reality they are often shaped by leadership choices, organizational culture, or incentive structures. A leadership team may tighten tolerance in response to reputational pressure or loosen it when uncertainty becomes difficult to ignore. Allowing these parameters to evolve endogenously would likely generate additional predictions about when organizations become more closed or more receptive to dissent.

Another limitation concerns the benefits of coordination. Some convergence in beliefs can be useful. Organizations require shared expectations to act coherently, and excessive disagreement can slow decision processes or create internal friction. The framework focuses on the informational cost of exclusion and therefore does not fully model the tradeoff between coordination and error under tail events. Similarly, the model does not specify the origin of the organizational anchor. Anchors may arise from leadership narratives, past experience, regulatory conventions, or prevailing market beliefs. Different sources may influence how strongly the anchor persists or how easily it shifts, but the mechanism of filtering around a focal belief remains the same. The mixture representation of regimes should also be understood as a tractable approximation of displacement events rather than a literal description of how structural change occurs. Real shocks often involve technological change, institutional transformation, or other forms of novelty that increase the value of diverse signals beyond what a simple mixture structure captures.

Several extensions follow naturally. One direction is to treat tolerance and conformity as endogenous choices shaped by organizational incentives or leadership strategies. Another is to model contributor turnover, particularly the exit of individuals who regularly produce dissenting or tail-relevant assessments. Persistent filtering could gradually deplete the population of such contributors, creating a feedback loop in which informational diversity erodes over time. A further extension would place the mechanism in a multi-firm setting. If institutions draw from the same talent pools, rely on similar models, or depend on common vendors, belief updating could become synchronized across firms. In that environment, recognition delays might occur simultaneously across institutions, raising

systemic concerns. Empirical work could support these directions by combining governance records, scenario documentation, hiring pipeline data, and time series of internal limit adjustments relative to external signals. Such evidence would help determine whether the patterns predicted by the model appear in real organizational decision processes.

8. Conclusions

Cohesion in financial organizations is not inherently stabilizing. When produced by exclusion-based filtering, it is variance depletion that concentrates beliefs around an anchor and can create illusory precision in normal times while degrading accuracy under displacement. The model formalizes how tolerance, conformity, hierarchy, and anchor updating jointly shape tail-regime bias, effective independence, and recognition delay, and it identifies observable diagnostics that can distinguish exclusion-based cohesion from integration-based cohesion before stress occurs.

Effective risk governance does not require eliminating filtering entirely. Coordination can require convergence, and not every dissenting signal is informative. The central implication is that governance systems should treat internal heterogeneity as a risk-management asset, and should monitor the mechanisms that determine whether heterogeneity is synthesized or excluded.

Funding: This research was funded by Willamette University, Atkinson Graduate School of Management.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A. Notation Table

Symbol	Definition	Range	Institutional Analog
θ_t	Latent risk state	$\mathbb{R}$	True exposure level
$\bar{s}_t$	Organizational anchor	$\mathbb{R}$	Dominant narrative/leadership prior
λ	Conformity parameter	$[0, 1)$	Career costs of divergence
κ	Report tolerance	$\mathbb{R}_+$	Formal or informal gatekeeping threshold
$\delta = \kappa / (1 - \lambda)$	Effective signal tolerance	$\mathbb{R}_+$	Combined exclusion radius
N_{eff}	Effective independence	$[1, N]$	Concentration of decision influence
α	Internal updating weight	$(0, 1)$	Responsiveness to internal filtered aggregate
η	External input weight	$[0, 1)$	Protected challenge function access
$\mathcal{A}_T$	Amplification ratio	$\mathbb{R}_+$	Tail-regime MSE relative to full-inclusion

Notes

- ¹ This paper uses the terms “anchor”, “anchor belief”, “dominant narrative”, “leadership prior”, and “institutional prior” interchangeably.

References

- Adams, R. B., & Ferreira, D. (2009). Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics*, 94(2), 291–309. [CrossRef]
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797–817. [CrossRef]
- Bénabou, R. (2013). Groupthink: Collective delusions in organizations and markets. *Review of Economic Studies*, 80(2), 429–462. [CrossRef]
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026. [CrossRef]
- Bikhchandani, S., & Sharma, S. (2001). Herd behavior in financial markets. *IMF Staff Papers*, 47(3), 279–310. [CrossRef]
- Csaszar, F. A., & Eggers, J. P. (2013). Organizational decision making: An information aggregation view. *Management Science*, 59(10), 2257–2277. [CrossRef]
- Economou, F., Gavriilidis, K., Goyal, A., & Kallinterakis, V. (2015). Herding dynamics in exchange groups: Evidence from Euronext. *Journal of International Financial Markets, Institutions and Money*, 34, 228–244. [CrossRef]
- Esser, J. K. (1998). Alive and well after 25 years: A review of groupthink research. *Organizational Behavior and Human Decision Processes*, 73(2–3), 116–141. [CrossRef] [PubMed]
- Federal Reserve Board of Governors. (2023). *Review of the Federal Reserve's supervision and regulation of Silicon Valley Bank*. Available online: <https://www.federalreserve.gov/publications/files/svb-review-20230428.pdf> (accessed on 15 February 2026).
- Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. Available online: <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/> (accessed on 15 February 2026).
- Janis, I. L. (1972). *Victims of groupthink: A psychological study of foreign-policy decisions and fiascos*. Houghton Mifflin.
- Ottaviani, M., & Sørensen, P. N. (2006). The strategy of professional forecasting. *Journal of Financial Economics*, 81(2), 441–466. [CrossRef]
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton University Press.
- Scharfstein, D. S., & Stein, J. C. (1990). Herd behavior and investment. *American Economic Review*, 80(3), 465–479.
- Valukas, A. R. (2010). Report of Anton R. Valukas, examiner. In *re Lehman Brothers Holdings Inc., et al., chapter 11 case no. 08-13555 (JMP)*. Bankr. S.D.N.Y.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.