

Volume 2, Issue 3

Research Article

Date of Submission: 06 July, 2026

Date of Acceptance: 03 August, 2026

Date of Publication: 11 August, 2026

When Markets Stop Listening: Stochastic Noise Intensity, the Generalized Signal Process, and the Design of Macroprudential Policy

Darius Joseph Hanson* 

Grand Canyon University, United States

Corresponding Author:

Darius Joseph Hanson, Grand Canyon University, United States.

Citation: Hanson, D. J. (2026). When Markets Stop Listening: Stochastic Noise Intensity, the Generalized Signal Process, and the Design of Macroprudential Policy. *Int Rev Bus Trade and Econ*, 2(3), 01-14.

JEL Classification: G32, G28, D82, C61, G01

Abstract

This paper extends the dynamic informational asymmetry framework established in Hanson (2025) by treating the noise intensity parameter η not as a fixed exogenous constant but as a time-varying stochastic process with its own governing dynamics [1]. The central limitation of prior models in this literature is that the informational environment itself is treated as stationary [2-5]. In financial crises, it manifestly is not. We introduce a stochastic differential equation for η_t that incorporates mean-reverting drift, stochastic volatility, and Poisson jumps, where jump intensity is endogenously linked to the credit risk state variable λ_t . We decompose η_t into a systematic component common to all institutions and an idiosyncratic component driven by firm-specific opacity, enabling a cross-sectional theory of noise contagion. Within this generalized framework, we formally derive the noise-dominated threshold η^* , the level above which the Kalman–Bucy belief-updating mechanism becomes statistically inert and managerial signaling fails regardless of the firm’s true solvency. This threshold generalizes the Strategic Exit Condition introduced in Hanson (2026) by adding a second, noise-triggered channel of institutional failure [6]. We then recast the dominant post-crisis regulatory instruments (mandatory disclosure requirements, the Liquidity Coverage Ratio, stress testing, and emergency liquidity facilities) as η -suppression mechanisms within a unified theoretical framework. Our primary policy contribution is a tiered dynamic activation rule that conditions regulatory intervention on the estimated distance between η_t and η^* , enabling preemptive macroprudential action before noise-dominated panics force the liquidation of solvent institutions. Calibration against market-based proxies demonstrates that Bear Stearns in March 2008 crossed the noise-dominated threshold primarily through the systematic component of η_t , while the Lehman Brothers failure in September 2008 reflected elevated idiosyncratic noise compounded by genuine deterioration in the latent value process V_t . These findings clarify the appropriate regulatory response for each failure type and carry direct implications for the design of forward-looking macroprudential oversight.

Keywords: Stochastic Noise Intensity, Informational Asymmetry, Macroprudential Policy, Liquidity Regulation, Systemic Risk, Signal Process, Bayesian Belief Updating

Introduction

The efficient transmission of information is the operational premise of modern finance. Fama (1970) establishes that competitive markets aggregate dispersed private information into prices, and Grossman and Stiglitz (1980) formalize the tension between information acquisition and market efficiency [7,8]. Yet as the financial crisis of 2007–2009 demonstrated with brutal clarity, the mechanism through which information is transmitted to prices is not itself invariant. There are environments in which the signal channel degrades: in which market participants, facing an ocean of structural complexity and rapidly evolving credit conditions, lose the capacity to distinguish fundamental value from noise. In those

environments, the theoretical apparatus of informationally efficient markets fails not because information asymmetry worsens, but because the transmission mechanism itself breaks down.

Hanson (2025) establishes a continuous-time stochastic framework in which firm value V_t and credit risk λ_t evolve as orthogonal stochastic processes, and investors observe a noisy aggregate signal Y_t [1]. The noise intensity parameter η in that signal process governs the efficacy of market-based belief updating. Hanson (2026) applies this framework to the collapse of Bear Stearns, demonstrating that the firm possessed substantial positive intrinsic value at the moment of its liquidation and was driven to a strategic exit not by insolvency but by an explosive drift in λ_t in a regime where η had risen to the point where standard managerial signaling was statistically ineffective [6]. The empirical companion characterizes this as a “noise-dominated” environment and validates the orthogonality of the latent state variables using the ex-post performance of the Maiden Lane LLC portfolio.

Both papers treat η as a fixed parameter of the informational environment. This is a convenient simplification that belies a critical empirical regularity: noise intensity is not constant. The complexity of structured financial products evolved over years. The opacity of mark-to-model Level 3 assets accumulated gradually as off-balance-sheet vehicles proliferated. The speed with which credit default swap spreads widened in the final week of March 2008 was not anticipated by any static measure of informational quality. The noise that overwhelmed the Bear Stearns signal process was itself a stochastic phenomenon; it evolved, it spiked, and it was partially correlated across institutions in ways that amplified rather than diversified systemic risk.

This paper addresses that gap by treating η_t as a dynamic state variable governed by its own stochastic differential equation. This single generalization has far-reaching consequences for both theory and policy. Theoretically, it transforms the noise-dominated regime from a static binary condition into a dynamic threshold-crossing problem. The question is no longer “is the market noise-dominated?” but rather “when will the noise process η_t next cross the critical threshold η^* , and how long will it remain above it?” Practically, this reframing opens the door to a genuinely anticipatory macroprudential policy framework. If η_t is observable—or at minimum estimable from market data—then regulators can monitor the distance between the current noise environment and the threshold at which market mechanisms fail, and intervene before that threshold is reached.

We make four principal contributions. First, we introduce a generalized stochastic process for η_t that captures both the gradual accumulation of opacity through financial innovation and the episodic, jump-driven spikes associated with crisis propagation. The jump intensity in this process is endogenously linked to the credit risk state variable λ_t , formalizing the intuition that liquidity stress and informational breakdown are not independent phenomena. Second, we decompose η_t into systematic and idiosyncratic components, providing a cross-sectional theory of noise contagion across institutions. The systematic component is driven by market-wide uncertainty and is common to all firms in a given funding market. The idiosyncratic component reflects firm-specific opacity, such as the ratio of Level 3 assets to tangible equity identified as a critical driver in Hanson (2026) [6]. Third, we derive the noise-dominated threshold η^* as a formal function of model parameters and show that crossing this threshold constitutes an independent channel of institutional failure, distinct from the solvency-based exit condition in Hanson (2025) [1]. This yields a generalized Strategic Exit Condition with two binding constraints rather than one. Fourth, we reinterpret post-crisis regulation within this framework and derive a dynamic tiered intervention rule that conditions the intensity of macroprudential action on the estimated proximity of η_t to η^* .

The paper proceeds as follows. Section 2 reviews the relevant literature on informational asymmetry, stochastic modeling in finance, and macroprudential policy design. Section 3 introduces the generalized stochastic noise process and its cross-sectional decomposition. Section 4 analyzes market dynamics under stochastic noise, deriving the noise-dominated threshold and generalizing the Strategic Exit Condition. Section 5 develops the theory of regulatory η -suppression, linking each major post-crisis instrument to the noise framework. Section 6 provides a calibration framework and applies the cross-sectional decomposition to the 2008 crisis. Section 7 proposes the dynamic macroprudential activation rule and discusses its implementation. Section 8 acknowledges limitations. Section 9 concludes.

Related Literature

This paper sits at the intersection of three literatures: the theory of informational asymmetry in financial markets, the stochastic modeling of financial dynamics, and the design of macroprudential regulatory policy. We briefly situate our contribution within each.

Informational Asymmetry in Corporate Finance

The foundational analysis of informational asymmetry in markets originates with Akerlof (1970), who demonstrates that quality uncertainty can collapse market mechanisms entirely when sellers possess private information that buyers cannot verify [1]. Leland and Pyle (1977) extend this logic to corporate finance, showing that entrepreneurs can credibly signal firm quality by retaining equity and accepting undiversified risk exposure [3]. Myers and Majluf (1984) establish that the informational advantage of managers over outside investors creates a pecking order of financing preferences, with internal funds preferred to debt and debt preferred to equity [4].

These foundational contributions share a critical structural feature: informational asymmetry is treated as a static endowment. The manager knows something investors do not, this knowledge is fixed at the time of the financing decision, and the problem is to design contracts or signals that reveal it. Agrawal and Rao (2015) advance the field by bifurcating informational asymmetry into two distinct factors (Value Informational Asymmetry (VIA) and Credit Risk Informational Asymmetry (CRIA)), demonstrating that many canonical models in corporate finance are special cases of this two-factor structure [5]. However, Agrawal and Rao largely retain the static assumptions of their predecessors.

Hanson (2025) relaxes the stationarity assumption by embedding both factors into a continuous-time stochastic framework, modeling V_t and λ_t as orthogonal Itô processes and specifying an observable noisy signal Y_t with constant noise intensity η [1]. Daley and Green (2012) study a related dynamic model in which news arrives as a Brownian diffusion process; their observation that the quality of news—not merely its valence—determines market responses anticipates our emphasis on the noise parameter [10]. The present paper extends Hanson (2025) by making η itself stochastic, which is, to our knowledge, a novel contribution to this literature [1].

Stochastic Modeling of Financial Dynamics

The use of stochastic differential equations in finance originates with Black and Scholes (1973) and Merton (1974) [10,11]. Merton's structural model of credit risk treats the firm's asset value as a geometric Brownian motion and derives default as the first-passage time of that process below the liability threshold. This framework implies a deterministic link between asset value and credit spreads, which the Bear Stearns case empirically refutes: the firm's credit spreads exploded while its asset value remained positive.

The limits of the Merton framework in crisis environments are extensively documented. Duffie and Singleton (1999) develop reduced-form models in which default intensity follows its own stochastic process, decoupled from the firm's fundamental value [12]. This is conceptually aligned with the orthogonality of V_t and λ_t in Hanson (2025), though derived from different assumptions. Longstaff and Schwartz (1995) incorporate interest rate risk into structural default models, and Collin-Dufresne, Goldstein, and Martin (2001) demonstrate empirically that credit spread changes are largely driven by factors unrelated to the standard structural determinants, lending credence to the view that λ_t is partially exogenous [1,13,14].

The introduction of jump processes into financial modeling is associated with Merton (1976) and Kou (2002), among others [15,16]. Our specification of the noise process η_t with Poisson jumps draws on this tradition. The regime-switching literature, initiated by Hamilton (1989) and applied to financial markets by Ang and Bekaert (2002) and others, provides an alternative framework for capturing the episodic nature of high-noise environments [17,18]. We favor the continuous-time SDE approach for consistency with Hanson (2025), but we discuss the regime-switching interpretation in Section 3 [1].

The Kalman–Bucy filter provides the continuous-time analogue of Bayesian updating in linear-Gaussian systems. Brendle and Carmona (2004) and Xiong (2001) develop filtering-based frameworks for financial markets with noisy signals [19,20]. Our characterization of the noise-dominated threshold η^* draws directly on the structure of the Kalman gain, which is the weight assigned to new observations in the updating equation.

Macroprudential Policy Design

The macroprudential policy literature has expanded substantially since the 2008 crisis. Borio (2003) provides an early framework distinguishing micro- and macroprudential approaches [21]. Brunnermeier and Sannikov (2014) develop a continuous-time model of the macroeconomy with financial intermediaries, demonstrating how amplification mechanisms can transform small shocks into systemic crises [22]. Their analysis of the “volatility paradox” (the observation that periods of low measured risk are associated with the buildup of systemic fragility) is directly relevant to the slow-moving accumulation of opacity (η_t drift) that precedes crisis spikes. Hanson, Kashyap, and Stein (2011) analyze the procyclicality of bank capital regulation and propose countercyclical capital buffers [23]. Adrian and Brunnermeier (2016) develop CoVaR as a measure of systemic risk and document the procyclical nature of financial institution leverage [24]. Gorton and Metrick (2012) provide the foundational empirical analysis of the “run on repo” mechanism, which Hanson (2026) interprets as a structural break in the λ_t process [6,25]. The present paper's contribution to this literature is the formal derivation of a noise-based criterion for regulatory intervention, complementing the existing toolkit of leverage-based, liquidity-based, and systemic-risk-based triggers.

The Liquidity Coverage Ratio (LCR) introduced by Basel III is analyzed by Kashyap, Stein, and Hanson (2010) and Banerjee and Mio (2018), among others [26,27]. The design of emergency liquidity facilities is discussed by Bagehot (1873) (“lend freely, at a penalty rate, against good collateral”) and its modern interpretations are reviewed by Rochet and Vives (2004) [28,29]. Our framework provides a theoretical foundation for the Bagehot criterion: institutions with positive V_t should receive liquidity support when $\eta_t > \eta^*$, because the private market has lost the informational capacity to distinguish them from genuinely insolvent counterparts.

The Generalized Noise Process

Limitations of Constant η

In Hanson (2025), the observable signal process is defined as [1]:

$$dY_t = \theta_v V_t dt + \theta_\lambda \lambda_t dt + \eta dB_t \quad (1)$$

where θ_v and θ_λ are loading parameters, B_t is a Brownian motion independent of the state variable innovations, and $\eta \in \mathbb{R}^+$ is the noise intensity. This specification captures the essential insight that investors observe an imperfect signal of the firm's true state, but it restricts the informational environment to be stationary in the sense that the signal-to-noise ratio is constant in expectation for given V_t and λ_t .

This restriction is empirically untenable in at least three dimensions. First, the opacity of financial instruments is not constant over time. The proliferation of structured finance products between 2003 and 2007 systematically increased the fraction of institutional assets that were "mark-to-model" rather than mark-to-market, gradually elevating the noise content of financial disclosures. Second, in crisis environments, noise spikes non-linearly and discontinuously. The repo haircut data documented by Gorton and Metrick (2012) shows near-zero haircuts on structured products prior to 2007, followed by exponential increases in early 2008 and effective 100% haircuts by March 2008, a pattern inconsistent with any stationary stochastic process [6]. Third, the noise environment is correlated across institutions. When one major counterparty faces a confidence crisis, the uncertainty about others in the same funding network rises simultaneously, creating systemic noise amplification that a firm-level constant cannot capture.

These three observations motivate the generalization that follows. We replace the constant η with a time-varying stochastic process η_t , and we specify the dynamics of that process to capture gradual drift, stochastic volatility, endogenous jumps, and cross-sectional correlation.

The Stochastic Noise Intensity Process

Let $\eta_t \geq 0$ denote the noise intensity at time t . We model η_t as a jump-diffusion process:

$$d\eta_t = \kappa_\eta(\bar{\eta} - \eta_t)dt + \sigma_\eta \eta_t dB_{\eta,t} + J_t dN_t \quad (2)$$

The components of this specification are as follows:

- $\kappa_\eta(\bar{\eta} - \eta_t)dt$: A mean-reverting drift toward the long-run noise level $\bar{\eta} > 0$, at speed $\kappa_\eta > 0$. In tranquil markets, η_t gravitates toward $\bar{\eta}$, which reflects the baseline level of opacity inherent in standard financial disclosures. The speed of reversion κ_η governs how quickly noise subsides after an acute episode.
- $\sigma_\eta \eta_t dB_{\eta,t}$: A proportional diffusion term, where $B_{\eta,t}$ is a standard Brownian motion. The proportionality to η_t ensures that noise fluctuations are larger when the noise level is already elevated—a variance-amplifying property consistent with the empirical behavior of financial uncertainty measures such as the VIX.
- $J_t dN_t$: A Poisson jump component, where N_t is a counting process with time-varying intensity $\rho(\lambda_t)$, and $J_t \sim \text{LogNormal}(\mu_j, \sigma^2)$ is the random jump size. We restrict $J_t > 0$, so all jumps are positive (noise spikes, not noise collapses). The key innovation is the specification of the jump intensity:

$$\rho(\lambda_t) = \rho_0 + \rho_1 \lambda_t \quad (3)$$

This formulation makes the probability of a noise spike an increasing function of the credit risk state variable. When λ_t is elevated, counterparties become more selective in their information processing, rumors spread more readily, and the probability of a discrete jump in η_t rises. This endogenous coupling between credit stress and informational breakdown is a central feature of the model and formalizes the observation in Hanson (2026) that the collapse of repo market confidence was not an exogenous event but was conditioned on the existing credit environment [6].

We assume throughout that the noise Brownian motion $B_{\eta,t}$ is independent of the state variable Brownian motions $W_{v,t}$ and $W_{\lambda,t}$ defined in Hanson (2025) [1]. This assumption preserves the structure of the original model while allowing η_t to fluctuate independently of the state variables. We relax this assumption in the discussion of feedback loops in Section 8.

For analytical tractability, we note that in the absence of jumps and with appropriate parameter restrictions, (2) reduces to the Cox–Ingersoll–Ross process (Cox, Ingersoll, and Ross, 1985), for which closed-form solutions to the distribution of η_t are available [30]. The full jump-diffusion specification is required to capture the acute spike dynamics observed in crisis periods.

Cross-Sectional Decomposition: Systematic and Idiosyncratic Noise

The specification in (2) models a single firm's noise intensity. To analyze cross-sectional dynamics and systemic contagion, we decompose $\eta_{i,t}$ for firm i as follows:

$$\eta_{i,t} = \eta_t^M + \eta_{i,t}^I \quad (4)$$

where η_t^M is the systematic (market-wide) noise component common to all institutions, and $\eta_{i,t}^I$ is the idiosyncratic noise component specific to firm i . We model each component separately

$$d\eta_t^M = \kappa^M(\bar{\eta}^M - \eta_t^M)dt + \sigma^M \eta_t^M dB_t^M + J_t^M dN_t^M \quad (5)$$

$$d\eta_{i,t}^I = \kappa^I(\bar{\eta}_i^I - \eta_{i,t}^I)dt + \sigma_i^I \eta_{i,t}^I dB_{i,t}^I \quad (6)$$

The systematic component η_t^M is driven by market-wide factors: the complexity of financial products prevalent in the market, the general opacity of the funding environment, and aggregate uncertainty. Its jump component N_t^M fires according to the market-wide credit environment, specified as $\rho(\Phi_t)$ where Φ_t is a market-level credit stress index. The idiosyncratic component $\eta_{i,t}^I$ follows a diffusion without jumps, reflecting the gradual evolution of firm-specific opacity driven by its balance sheet composition. We omit idiosyncratic jumps for parsimony, though they could be incorporated for institutions facing sudden asset reclassification.

The long-run mean $\bar{\eta}_i^I$ captures firm-specific structural opacity. We propose the following empirical proxy:

$$\bar{\eta}_i^I = \alpha_0 + \alpha_1 \left(\frac{L3_i}{TCE_i} \right) + \alpha_2 \left(\frac{OTC_i}{TA_i} \right) \quad (7)$$

where $L3/TCE_i$ is the ratio of Level 3 assets to tangible common equity and OTC_i/TA_i is the ratio of over-the-counter derivative notional exposure to total assets. As demonstrated in Hanson (2026), Bear Stearns' Level 3 to TCE ratio of approximately 252% at year-end 2007 was a primary driver of its elevated idiosyncratic noise level [6]. This ratio provides a balance sheet-based predictor of the firm's structural informational opacity prior to any crisis event.

The decomposition in (4)–(6) has two important implications for regulatory policy. First, it separates the sources of noise into market-level and firm-level drivers, implying distinct policy responses: macroprudential tools should target η_t^M , while microprudential disclosure requirements should target $\eta_{i,t}^I$. Second, it formalizes noise contagion: when η_t^M spikes due to a crisis event, all institutions simultaneously experience an increase in their total noise level $\eta_{i,t}^I$, regardless of their idiosyncratic opacity. Institutions with high $\bar{\eta}_i^I$ are therefore “pre-positioned” at a higher total noise level and require a smaller increment from the systematic component to cross the noise-dominated threshold.

Market Dynamics Under Stochastic Noise The Modified Signal Process

Replacing the constant η in the original signal process with the stochastic $\eta_{i,t}$, the observable signal for firm i becomes:

$$dY_{i,t} = \theta_v V_{i,t} dt + \theta_\lambda \lambda_{i,t} dt + \eta_{i,t} dB_{i,t} \quad (8)$$

This is a nonlinear filtering problem. The investor's posterior distributions of $V_{i,t}$ and $\lambda_{i,t}$ conditional on the observation history $\{Y_{i,s}, s \leq t\}$ are no longer characterized by the standard linear Kalman–Bucy filter, because the observation noise variance $\eta_{i,t}^2$ is itself stochastic and unobserved. This is a state-space model with stochastic observation noise, a class studied by, among others, Liu and West (2001) in the particle filtering literature [31].

For analytical purposes, we proceed with the assumption that investors condition on a filtration generated by both $Y_{i,t}$ and a market-observable proxy $\hat{\eta}_t$ for the noise level (developed explicitly in Section 6). Under this assumption, the conditional filtering equations take the form of the standard Kalman–Bucy system with a time-varying noise variance $\eta_{i,t}^2$:

$$d\hat{V}_{i,t} = \alpha_v \hat{V}_{i,t} dt + K_{v,i,t} (dY_{i,t} - \theta_v \hat{V}_{i,t} dt - \theta_\lambda \hat{\lambda}_{i,t} dt) \quad (9)$$

$$K_{v,i,t} = P_{v,i,t} \frac{\theta_v}{\eta_{i,t}^2} \quad (10)$$

where $\hat{V}_{i,t}$ denotes the investor's posterior mean estimate of firm value, $P_{v,i,t}$ is the posterior variance satisfying a Riccati-type ordinary differential equation, and $K_{v,i,t}$ is the Kalman gain: the weight assigned to the innovation (the new information content of the observed signal increment). An analogous equation holds for $\hat{\lambda}_{i,t}$ with gain $K_{\lambda,i,t} = P_{\lambda,i,t} \theta_\lambda / \eta_{i,t}^2$.

The Noise-Dominated Threshold

The Kalman gain $K_{v,i,t}$ in equation (10) is the central object of interest for our policy analysis. Specifically, as $\eta_{i,t} \rightarrow \infty$, $K_{v,i,t} \rightarrow 0$. When $K_{v,i,t} = 0$, the updating equation (9) reduces to $d\hat{V}_{i,t} = \alpha_v \hat{V}_{i,t} dt$: the investor's belief evolves purely according to prior dynamics, with zero weight on new market observations. In this regime, the signal process $Y_{i,t}$ ceases to perform any price discovery function. Managerial signaling—the central mechanism by which managers with positive private information can stabilize market perceptions in the Hanson (2025) framework—is rendered statistically invisible [1].

We formalize this observation by defining the noise-dominated threshold. For a given minimum effective gain $\varepsilon > 0$ below which belief updating is considered “inert,” the noise-dominated threshold η^* is defined implicitly by:

$$K_{v,i,t} = \varepsilon \iff \eta_{i,t}^* = \sqrt{\frac{P_{v,i,t}\theta_v}{\varepsilon}} \quad (11)$$

Two features of (11) are important. First, η^* is itself time-varying, because it depends on the posterior variance $P_{v,i,t}$. As the investor's uncertainty about firm value increases ($P_{v,i,t}$ rises), so does the threshold: a noisier prior makes the investor more reliant on new signals, so a higher noise level is required to render those signals inert. Second, η^* depends on the loading parameter θ_v : assets with less transparent fundamental values (lower θ_v) require less noise to push the signal into the noise-dominated regime.

Combining (11) with the decomposition in (4), the noise-dominated condition for firm i is:

$$\eta_t^M + \eta_{i,t}^I > \sqrt{\frac{P_{v,i,t}\theta_v}{\varepsilon}} \quad (12)$$

This condition reveals the asymmetric vulnerability of institutions with high idiosyncratic noise. When the systematic component η_t^M spikes during a market-wide crisis, institutions with elevated $\eta_{i,t}^I$ breach the noise-dominated threshold first. They are, in a precise mathematical sense, more susceptible to noise-triggered failure than their peers—not because their fundamentals are worse, but because their informational structure is more opaque.

The Generalized Strategic Exit Condition

Hanson (2025) derives the optimal manager action $a_t^* = V_t$ from the first-order condition of the instantaneous utility function $u(a_t, V_t, \lambda_t) = a_t V_t - (1/2)a_t^2 - \gamma\lambda_t$. Hanson (2026) shows that the utility u falls below the manager's reservation utility (zero) and triggers a strategic exit when [1,6]:

$$\lambda_t > \frac{V_t^2}{2\gamma} \quad (13)$$

This is the original Strategic Exit Condition, driven entirely by the credit risk process λ_t overwhelming the value-based return. The introduction of stochastic η_t opens a second, independent channel through which a firm can be forced into exit even when condition (13) does not bind. Specifically, when $\eta_t > \eta^*$, the market's Kalman gain approaches zero, so investors can no longer verify the manager's privately observed V_t . In the absence of credible updating, counterparties act on the basis of their stale priors. If those priors are sufficiently pessimistic (as they were in the case of Bear Stearns after the collapse of its hedge funds in 2007), the effective λ_t as perceived by the market diverges from the true λ_t .

We formalize this with the following definition. Let $\hat{\lambda}_t$ denote the market's posterior estimate of the credit risk process given the noise-dominated signal. Define the perceived

$$\hat{u} = a_t \hat{V}_t - \frac{1}{2}a_t^2 - \gamma\hat{\lambda}_t \quad (14)$$

When $\eta_t > \eta^*$, $\hat{V}_t$ and $\hat{\lambda}_t$ evolve according to prior dynamics with zero weight on new observations. Over any period of length Δt in the noise-dominated regime:

$$\hat{V}_t \rightarrow E[V_t | \mathcal{F}_{T_\eta}], \quad \hat{\lambda}_t \rightarrow E[\lambda_t | \mathcal{F}_{T_\eta}] \quad (15)$$

where $\mathcal{F}_{T_\eta}$ is the investor's information set at the moment the noise-dominated regime was entered (time T_η). If $E[\lambda_t | \mathcal{F}_{T_\eta}]$ is large relative to $E[V_t | \mathcal{F}_{T_\eta}]$ due to the accumulated pessimism of the pre-crisis period, the perceived utility $\hat{u}$ can be negative even when the true utility u is positive. This constitutes what we term a noise-triggered strategic exit: a rational liquidation decision driven not by fundamental deterioration but by the informational breakdown itself.

Definition 1 (Generalized Strategic Exit Condition): A firm with utility function $u(a_t, V_t, \lambda_t)$ exits at the first time τ satisfying either of the following conditions:

- **Condition A (Solvency-Driven):** $\lambda_t > V_t^2 / (2\gamma)$
- **Condition B (Noise-Driven):** $\eta_t > \eta^*$ and $E[\hat{u}|\hat{\lambda}_t] < 0$

The exit time is $\tau = \min(\tau_A, \tau_B)$, where τ_A and τ_B are the first times Condition A and Condition B, respectively, bind. The firm need not be insolvent at τ . Hanson (2026) establishes that Bear Stearns' exit was driven by τ_B , not τ_A [6].

This generalization has a critical policy implication: the appropriate regulatory response to Condition A exit is recapitalization or orderly insolvency resolution, while the appropriate response to Condition B exit is η -suppression (actions that reduce η_t below η^* before the exit time τ_B binds). The failure to distinguish between these two conditions in 2008 explains much of the regulatory inconsistency observed across institutions during that period.

Regulatory Design as η -Suppression

We now reinterpret the major post-crisis regulatory instruments within the stochastic noise framework. Each instrument operates on one of the key model parameters ($\eta_{i,t}$, $\lambda_{i,t}$, or the threshold η^*) and their effectiveness in preventing noise-triggered failures can be evaluated accordingly.

Mandatory Disclosure as Direct η -Reduction

Mandatory disclosure requirements oblige firms to release verifiable information to the market at specified intervals or upon specified triggers. Within our framework, a mandatory disclosure event at time t^D injects a signal Y^D with lower noise than the continuous market signal. Define the disclosure signal process:

$$dY_t^D = \theta_v^D V_t dt + \theta_\lambda^D \lambda_t dt + \eta^D dB_t^D \quad (16)$$

where $\eta^D < \eta_{i,t}$ in the market signal (16 incorporates independently verified information) and B_t^D is independent of the market noise process. The composite signal available to investors when both channels operate simultaneously has effective noise:

$$\eta^{\text{eff}} = \frac{\eta_{i,t} \cdot \eta^D}{\sqrt{\eta_{i,t}^2 + (\eta^D)^2}} < \min(\eta_{i,t}, \eta^D) \quad (17)$$

Equation (17) demonstrates that mandatory disclosure always strictly reduces the effective noise below the market noise level. When $\eta_{i,t}$ is large (a high-noise market regime), $\eta^{\text{eff}} \approx \eta^D$, meaning the disclosure signal dominates. This provides the theoretical basis for the intuition that transparency requirements are most valuable precisely when markets are most opaque.

The policy implication is a frequency and granularity prescription: mandatory disclosure intervals should be shorter and disclosure content should be more granular (higher θ_v^D , higher θ_λ^D) for institutions with high $\eta_{i,t}^I$, as proxied by the Level 3 / TCE ratio defined in (7). Equivalently, existing fair value accounting standards under SFAS 157 should impose mandatory Level 3 to Level 2 reclassification when the ratio L3/TCE exceeds a threshold χ^* , triggering an automatic reduction in $\eta_{i,t}^I$. We return to the calibration of χ^* in Section 6.

The Liquidity Coverage Ratio as λ -Drift Suppression

The Liquidity Coverage Ratio (LCR) introduced under Basel III requires covered institutions to hold a buffer of High-Quality Liquid Assets (HQLA) sufficient to cover projected net cash outflows over a 30-day stress scenario. Within the Hanson (2025) framework, this instrument operates primarily on the credit risk process $\lambda_{i,t}$ rather than directly on $\eta_{i,t}$ [1].

Recall that the credit risk process is specified in Hanson (2025) as $d\lambda_t = \alpha_\lambda \lambda_t dt + \sigma_\lambda dW_{\lambda,t}$ [1]. Bear Stearns funded \$102 billion in daily repo, so the effective volatility σ_λ was extremely high. The LCR addresses this by requiring a mandatory HQLA buffer $H_{i,t}$ that absorbs the first tranche of outflows, effectively replacing the uncapped drift process with a constrained version. Define the LCR-modified credit risk process:

$$d\lambda_{i,t}^{\text{LCR}} = d\lambda_{i,t} \cdot \mathbf{1}(\lambda_{i,t} > H_{i,t}) \quad (18)$$

where $\mathbf{1}(\cdot)$ is the indicator function. The LCR requirement imposes $H_{i,t} \geq \bar{H}$, a minimum HQLA level. As long as $\lambda_{i,t} \leq H_{i,t}$, outflows are absorbed from the buffer and $d\lambda_{i,t}^{\text{LCR}} = 0$. Only when the buffer is exhausted does $\lambda_{i,t}$ begin its stochastic drift.

The LCR therefore operates by extending the time before the credit risk process $\lambda_{i,t}$ reaches the exit threshold in condition (13). Formally, it increases the first-passage time τ_A . Crucially, it also indirectly delays τ_B : because the jump intensity $\rho(\lambda_{i,t})$ in equation (3) is an increasing function of $\lambda_{i,t}$, suppressing $\lambda_{i,t}$ drift reduces the probability of a noise spike that triggers Condition B. This indirect channel (the LCR reducing noise-triggered exits by keeping $\lambda_{i,t}$ low, which keeps noise jump intensity low) is the theoretically novel contribution of the present analysis.

The LCR is therefore doubly beneficial in our framework: it directly delays Condition A failure and indirectly delays Condition B failure through the $\rho(\lambda_{i,t})$ linkage. This provides a theoretical grounding for the Basel Committee's emphasis on liquidity regulation as a systemic stability tool, beyond the firm-level solvency rationale that dominates existing justifications.

Stress Testing as Regulatory η -Verification

Mandatory stress tests, operationalized in the United States through the Comprehensive Capital Analysis and Review (CCAR) and Dodd-Frank Act Stress Tests (DFAST), require covered institutions to project their capital positions under severely adverse economic scenarios, with results verified by regulatory examiners. Within our framework, this process constitutes a regulatory injection of a verification signal $Y_{i,t}^{ST}$ with substantially lower noise than the private market signal:

$$dY_{i,t}^{ST} = \theta_v^S V_{i,t} dt + \theta_\lambda^S \lambda_{i,t} dt + \eta^S dB_{i,t}^S \quad (19)$$

where $\eta^S \ll \eta_{i,t}$ because regulators have access to non-public data, independent valuation methodologies, and cross-institutional benchmarking capacity unavailable to private market participants. From (17), combining the stress test signal with the market signal yields:

$$\eta_S^{\text{eff}} = \frac{\eta_{i,t} \cdot \eta^S}{\sqrt{\eta_{i,t}^2 + (\eta^S)^2}} \approx \eta^S \quad \text{when } \eta^S \ll \eta_{i,t} \quad (20)$$

In the high-noise market environment, the stress test signal essentially replaces the market signal as the operative information source. This has two implications. First, the timing of stress test disclosure relative to the noise cycle matters. A stress test result disclosed when η_t^M is already elevated will fail to suppress $\eta_{i,t}$ below η^* , because by the time the market processes the regulatory signal, the noise-dominated regime may have already triggered exit. This is precisely what occurred with CEO Schwartz's CNBC appearance on March 12, 2008: a credible signal in a low-noise environment would have been stabilizing; in the high-noise regime of that week, it was statistically inert.

Second, the frequency and design of stress tests should be countercyclical with respect to η_t^M . In low-noise environments (normal markets), annual stress tests with delayed disclosure may be adequate. As η_t^M rises, stress test frequency should increase and disclosure should accelerate, maximizing the period during which $\eta_S^{\text{eff}} < \eta^*$. We formalize this in the dynamic activation rule in Section 7.

Emergency Liquidity Facilities as Threshold-Shifting Interventions

When neither disclosure nor liquidity requirements prevent $\eta_{i,t}$ from breaching η^* , the remaining regulatory option is direct intervention in the credit environment to reduce $\lambda_{i,t}$. The Federal Reserve's extension of an emergency lending facility to Bear Stearns on March 14, 2008 via JPMorgan, followed by the creation of the Primary Dealer Credit Facility (PDCF), constitutes exactly this type of intervention in our framework.

Formally, the emergency facility imposes a cap $\bar{\lambda}$ on the credit risk process: by guaranteeing to provide unlimited liquidity at a penalty rate r_p , the facility ensures that any firm meeting the Bagehot solvency criterion ($V_{i,t} > 0$) can access funding regardless of private market sentiment. In the model, this is equivalent to truncating the $\lambda_{i,t}$ process at $\bar{\lambda}$:

$$\lambda_{i,t}^{\text{ELA}} = \min(\lambda_{i,t}, \bar{\lambda}) \quad (21)$$

This truncation has a direct effect on $\eta_{i,t}$ through the jump intensity linkage in (3): when $\lambda_{i,t}^{\text{ELA}}$ is bounded by $\bar{\lambda}$, the Poisson intensity $\rho(\lambda_{i,t}^{\text{ELA}})$ is bounded by $\rho_0 + \rho_1 \bar{\lambda}$, which suppresses further noise jumps. The emergency facility therefore not only prevents Condition A failure directly, but reduces the probability of subsequent Condition B failures by removing the $\lambda_{i,t} \rightarrow \eta_t$ feedback loop. This is the formal mechanism underlying the observation in Hanson (2026) that the subsequent renegotiation of the Bear Stearns sale price from \$2 to \$10 per share was possible once the liquidity panic (λ_i) had subsided: with $\lambda_{i,t}$ bounded, the noise-dominated regime collapsed, and the signal process $Y_{i,t}$ was again able to convey information about $V_{i,t}$ [6].

The appropriate penalty rate r_p in the Bagehot framework maps naturally to the model. The facility lender (central bank) observes that $\eta_{i,t} > \eta^*$ and that private market creditors cannot distinguish Condition A from Condition B failure. The

penalty rate compensates for the option value of lending to a potentially insolvent institution in a noisy environment. Within the model, the fair penalty rate is a function of the probability of Condition A failure conditional on the information available in the managerial filtration $\mathcal{F}_t^M$ (which the central bank can approximate through regulatory examination), scaled by the expected loss given default.

Calibration Framework and Cross-Sectional Evidence

Market-Based Proxies for η_t

The empirical implementation of the dynamic noise framework requires observable proxies for η_t^M and $\eta_{i,t}^I$. We propose the following two-component estimator:

For the systematic component η_t^M , we draw on market-wide uncertainty measures:

- VIX_t : The CBOE Volatility Index measures implied volatility from S&P 500 options and reflects broad market uncertainty. Elevated VIX corresponds to elevated η_t^M through the channel of heightened investor ambiguity about macroeconomic fundamentals.
- TED_t : The TED spread (3-month Libor minus 3-month T-bill rate) measures the premium charged for interbank lending relative to risk-free borrowing. This spread reflects the credit uncertainty of financial counterparties and maps directly to the noise in the in-terbank signal. Rising TED spreads elevate the effective η_t^M in the wholesale funding market.
- $MOVE_t$: The MOVE index measures implied volatility in the U.S. Treasury market, capturing rate uncertainty that complicates the valuation of fixed-income assets and therefore raises the noise content of mark-to-model estimates for institutions with large fixed-income books.

We propose the following composite systematic noise proxy:

$$\hat{\eta}_t^M = \beta_0 + \beta_1 \cdot VIX_t + \beta_2 \cdot TED_t + \beta_3 \cdot MOVE_t + \beta_4 \cdot \Sigma CDS_vol_t \quad (22)$$

where ΣCDS_vol_t is the cross-sectional standard deviation of 5-year CDS spread changes across major financial intermediaries. This last term captures the dispersion in market assessments of counterparty credit quality: when investors disagree sharply about which institutions are safe, the effective noise in the aggregate signal is high. The β coefficients can be estimated via panel regression on measures of signal extraction quality, such as the R^2 of analyst earnings forecast regressions or the bid-ask spread of index CDS products.

For the idiosyncratic component $\eta_{i,t}^I$, the primary balance sheet driver is the Level 3 / TCE ratio, as specified in (7). Additional determinants include:

- $OTC_{i,t} / TA_{i,t}$: Ratio of over-the-counter derivative notional to total assets. OTC positions are bilateral and not exchange-cleared, with valuation highly sensitive to model assumptions. Higher OTC concentration raises $\eta_{i,t}^I$.
- $Repo_{i,t} / TL_{i,t}$: Ratio of overnight repo funding to total liabilities. Heavy reliance on overnight repo creates informational opacity about the stability of the firm's funding base, because counterparties can exit positions overnight before any public signal of stress is released. This ratio was approximately 26% for Bear Stearns at year-end 2007.

Cross-Institutional Calibration: 2008 Failures

We apply the cross-sectional decomposition to contrast two canonical 2008 institutional failures: Bear Stearns (March 2008) and Lehman Brothers (September 2008) [32]. The Hanson (2026) analysis of Bear Stearns establishes that the firm's collapse was a Condition B failure: $V_t > 0$ throughout, λ_t crossed the exit threshold due to a noise-triggered run, and the Maiden Lane LLC portfolio subsequently generated a net gain of \$2.5 billion for the public, validating positive intrinsic value [6]. We contrast this with the Lehman case, which the limitations discussion in Hanson (2026) explicitly distinguishes as a different failure type [6].

The key structural difference between the two institutions is the composition of idiosyncratic noise $\eta_{i,t}^I$. Bear Stearns held a Level 3 / TCE ratio of approximately 252% at year-end 2007, driven primarily by CDO and MBS positions. Lehman Brothers held analogous ratios for its commercial real estate book, but the nature of the opacity differed: Lehman's CRE positions were not merely mark-to-model but were also subject to genuine deterioration in the underlying collateral values, meaning that the signal obscured by idiosyncratic noise was not a healthy V_t but an already-declining one.

In the framework of the Generalized Strategic Exit Condition, Bear Stearns was a pure Condition B failure: $\tau_B < \tau_A$. Lehman represents a mixed case: Condition A failure was approaching independently, and Condition B failure, driven by the same market-wide noise spike of September 2008, compounded and accelerated it. The implication is that the Bear Stearns collapse was eminently preventable through η -suppression, while the Lehman collapse required both η -suppression and genuine recapitalization.

The η -Gap at Bear Stearns: March 10–14, 2008

Using the proxy specification in (22), we can reconstruct the estimated η -gap $\Gamma_t = \eta^* - \hat{\eta}_t$ for Bear Stearns during its final week. The analysis draws on data documented in Hanson (2026) [6] and the primary sources cited therein.

Metric	Bear Stearns (Mar 2008)	Lehman Brothers (Sep 2008)	Policy Implication
Level 3 / TCE Ratio	~252%	~n.a. (CRE-driven)	Idiosyncratic η driver
Repo / Total Liabilities	~26%	~25%	High $\rho(\lambda_t)$ sensitivity
VIX at Failure Date	~32	~33	Similar η_t^M spike
TED Spread at Failure	~150 bps	~175 bps	Systematic stress
$V_t > 0$ at Failure?	Yes (Maiden Lane +\$2.5B)	Contested losses)	(CRE Exit condition type
Primary Driver	Exit Condition B (η -triggered)	Mixed A + B	Intervention type
Appropriate response	Re- η -Suppression (LCR/ELA)	Recapitalization + η support	Regulatory pre-scription

Table 1: Cross-Institutional Noise Decomposition, 2008

On March 10, the VIX stood at approximately 27, the TED spread was approximately 150 basis points, and Bear Stearns' 5-year CDS spreads were approximately 100 basis points with moderate volatility. These figures imply a moderate $\hat{\eta}_t^M$, well below our estimated η^* for a firm with Bear's loading parameters. The η -gap was positive: the market was functioning, and the signal process, while noisy, was providing meaningful information.

By March 13, the situation had deteriorated dramatically. CDS spreads had widened to over 700 basis points in 24 hours, the TED spread was spiking, and the cross-sectional standard deviation of bank CDS spread changes ($\Sigma_{\text{CDS_vol}}$) was at its highest level since the onset of the subprime crisis. Our estimated $\hat{\eta}_t$ had crossed the estimated η^* by the close of business on March 13, 2008, precisely when the FCIC (2011) documents the liquidity pool collapsing from \$12.5 billion to \$2 billion in a single day [33]. This is the moment at which the noise-dominated regime bound, and the remaining signaling capacity of the management team became statistically negligible. The \$12.9 billion emergency loan extended on March 14 was consumed within hours precisely because η_t was still above η^* : even with a government backstop, counterparties were unable to update their posterior beliefs about Bear's solvency quickly enough to halt the exit process already in motion.

Date	Liquidity Pool	CDS Spread	Est. $\hat{\eta}_t$	Est. η^*	η -Gap Γ_t	Regime
Mar 10	\$18.1B	~100 bps	Moderate	Fixed	+	Functional
Mar 11	~\$15.0B	~200 bps	Elevated	Fixed	Narrowing	Watch
Mar 12	\$12.5B	~400 bps	High	Fixed	Near-zero	Stress
Mar 13	\$2.0B	~700 bps	Critical	Fixed	<0	Noise-Dominated
Mar 14	~\$0 (ex-ELA)	~1000 bps	Extreme	Fixed	$\ll 0$	Exit Triggered

Table 2: Estimated η -Gap Trajectory, Bear Stearns (March 2008)

The Dynamic Macroprudential Activation Rule

The η -Gap as a Policy Trigger

The core policy innovation of this paper is the proposal that macroprudential regulators should continuously monitor the estimated η -gap:

$$\Gamma_{i,t} = \eta_{i,t}^* - \hat{\eta}_{i,t} \quad (23)$$

where $\eta_{i,t}^*$ is the estimated noise-dominated threshold for firm i at time t (calculated using (11) with regulatory estimates of the posterior variance $P_{v,i,t}$ and loading parameters) and $\hat{\eta}_{i,t}$ is the composite noise proxy constructed from (22) and (7). A positive $\Gamma_{i,t}$ indicates that the firm's signal process is functioning below the noise-dominated threshold. A negative or near-zero $\Gamma_{i,t}$ is a warning signal that the informational mechanism is approaching or has crossed the noise-dominated regime.

The monitoring of $\Gamma_{i,t}$ across the universe of systemically important financial institutions (SIFIs) provides regulators with a forward-looking indicator of market mechanism failure, rather than the backward-looking capital adequacy ratios that constituted the primary regulatory metric in 2008. The SEC's capital adequacy monitoring of Bear Stearns correctly identified solvency (it found no breach of capital requirements as late as March 11, 2008), but it provided no warning of

the approaching noise-triggered exit, because the noise parameter was not part of the regulatory monitoring framework.

Tiered Intervention Architecture

We propose a tiered intervention architecture conditioned on $\Gamma_{i,t}$, analogous to the prompt corrective action (PCA) framework for solvency but operating on the noise dimension:

Tier	η -Gap Condition	Label	Required Action	Authority
0	$\Gamma_{i,t} > 2\sigma_\eta$	Normal	Standard quarterly reporting	Microprudential
1	$\sigma_\eta < \Gamma_{i,t} \leq 2\sigma_\eta$	Watch	Enhanced weekly disclosure; stress test trigger	Microprudential
2	$0 < \Gamma_{i,t} \leq \sigma_\eta$	Stress	Mandatory public disclosure; HQLA draw-down audit; regulatory verification signal	Joint micro/macro
3	$\Gamma_{i,t} \leq 0$	Crisis	Emergency liquidity facility activation; OLA pre-positioning; FDIC receivership preparation	Macroprudential / FSOC

Table 3: Dynamic η -Monitor Intervention Tiers

The threshold parameters ($2\sigma_\eta$, σ_η) are expressed in units of the estimated standard deviation of the noise process, exploiting the parameterization of (2). This renders the intervention thresholds invariant to the level of η^* across different institutions and market environments, facilitating consistent regulatory application.

Two features of this architecture deserve emphasis. First, Tier 1 activation triggers an accelerated stress test, consistent with the countercyclical stress testing frequency prescription derived in Section 5.3. The regulatory verification signal $Y_{i,t}^{\text{ST}}$ injected at Tier 1 extends the window before the noise-dominated threshold is breached, buying time for the other stabilizing mechanisms to operate. Second, Tier 3 activation involves pre-positioning of Orderly Liquidation Authority resources before the institutional failure has actually occurred. This is the critical departure from the 2008 framework: rather than improvising a Bear Stearns resolution in 96 hours, a Tier 3 designation would initiate the FDIC's OLA pre-positioning process while the institution was still technically solvent, enabling an orderly stabilization rather than a fire sale.

The Systemic η -Monitor: Cross-Institutional Aggregation

Beyond the firm-level $\Gamma_{i,t}$, we propose a systemic noise index Ξ_t defined as the value-weighted harmonic mean of $\Gamma_{i,t}$ across all SIFIs:

$$\Xi_t = \left(\sum_i \frac{w_{i,t}}{\Gamma_{i,t}} \right)^{-1} \quad (24)$$

where $w_{i,t}$ is the systemic importance weight of institution i (proxied by its share of total SIFI assets). The harmonic mean structure is deliberate: it is dominated by institutions with the smallest $\Gamma_{i,t}$, ensuring that a single large institution approaching the noise-dominated threshold pulls the systemic index sharply downward. When Ξ_t falls below a critical threshold Ξ^* , the regulatory response should shift from institution-specific to market-wide, targeting η_M directly through broad-based emergency liquidity facilities such as the PDCF or the Term Asset-Backed Securities Loan Facility (TALF).

The systemic index Ξ_t provides a continuous, real-time analogue to the Financial Stability Oversight Council's designation process, which operates on an annual or at best quarterly basis. Mandating public disclosure of Ξ_t , analogously to the FSOC's annual report, would itself constitute an η -suppression mechanism, by aligning private market participants' expectations about the systemic noise environment with the regulatory assessment.

Limitations

The framework developed in this paper extends Hanson (2025) and Hanson (2026) in important directions, but it introduces new assumptions and abstractions that warrant explicit acknowledgment [1,6].

Observability of η_t and the Proxy Problem

The policy framework in Section 7 relies on the continuous estimation of $\hat{\eta}_{i,t}$ from market-based proxies. The mapping from observables (VIX, TED spread, CDS volatility, Level 3 ratios) to η_t is estimated, not exact. The β coefficients in (22) are themselves time-varying in a manner we have not modeled, and the loading parameters θ_v and θ_λ required to compute η^* from (11) are latent and must be estimated from historical data. An intervention framework based on imprecise $\Gamma_{i,t}$ estimates is subject to both Type I errors (premature intervention in functioning institutions) and Type II errors (failure to intervene in genuinely noise-dominated firms). The optimal tuning of the intervention thresholds requires a calibration exercise using historical crisis data that is beyond the scope of the present theoretical paper.

Linearity and Proportionality Assumptions The utility function $u(a_t, V_t, \lambda_t) = a_t V_t - (1/2)a_t^2 - \gamma \lambda_t$ is linear in both V_t and λ_t . This specification generates the clean result $a_t^* = V_t$, but it abstracts from important nonlinearities in managerial decision-making during distress. In practice, the utility of continued operation likely falls nonlinearly as λ_t rises, because reputational damage, regulatory scrutiny, and operational disruption accelerate at higher levels of credit stress. A more general utility specification would allow for convexity in λ_t and could generate richer predictions about the timing of exit relative to the noise-dominated threshold.

Feedback Between η_t and the State Variables

We assume throughout that the noise process $B\eta_t$ is independent of the state variable innovations $W_{v,t}$ and $W_{\lambda,t}$. While the endogenous coupling of jump intensity $\rho(\lambda_t)$ partially relaxes this independence, the diffusion components remain uncorrelated. In practice, elevated η_t likely feeds back into the drift and volatility of V_t and λ_t through at least two channels. First, when markets are noise-dominated, capital allocation is distorted: high-quality firms face elevated funding costs and may forgo positive-NPV investments, gradually reducing V_t . Second, sustained noise-dominated regimes can induce real economic contractions that elevate λ_t even for fundamentally solvent institutions. A full equilibrium model that endogenizes these feedback loops is a priority for future research.

Single-Firm Focus and Network Externalities

The cross-sectional decomposition in Section 3.3 captures noise contagion through the common systematic component η_t^M , but it does not model direct network externalities between specific pairs of institutions. In the Bear Stearns case, the failure of its internal hedge funds in June 2007 raised idiosyncratic noise not just for Bear but for other prime brokerage clients, repo counterparties, and CDO counterparties with similar exposures. A network model in the spirit of Elliott, Golub, and Jackson (2014) or Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) that propagates $\eta_{i,t}^I$ shocks through the financial network would provide a richer description of systemic noise amplification [34,35].

Conclusion

This paper makes the case for treating the noise intensity of financial markets as a dynamic object that is, itself, a proper subject of regulatory attention. Building on the theoretical framework established in Hanson (2025) and the empirical validation provided in Hanson (2026), we demonstrate that the constant noise parameter η of the original signal process is both theoretically restrictive and empirically inadequate for characterizing the informational environment of crisis markets [1,6]. The generalization introduced here (treating η_t as a jump-diffusion process with mean-reverting drift and endogenous jump intensity) resolves this inadequacy while preserving the analytical tractability of the underlying stochastic framework.

The primary theoretical contribution is the derivation of the noise-dominated threshold η^* and the Generalized Strategic Exit Condition, which establishes that institutions can be forced into liquidation by two independent mechanisms: the classical solvency-driven failure (Condition A) in which credit risk overwhelms value, and the novel noise-driven failure (Condition B) in which the informational mechanism through which value can be signaled breaks down entirely. The identification of these two distinct failure modes is not merely a theoretical refinement. It has direct implications for the design of regulatory interventions: Condition A failures require recapitalization or orderly resolution, while Condition B failures require η -suppression through disclosure, liquidity support, and regulatory verification.

The cross-sectional decomposition of $\eta_{i,t}$ into systematic and idiosyncratic components provides a novel theoretical basis for the intuition, long present in the macroprudential literature but poorly formalized, that financial opacity is a systemic risk. Institutions with high idiosyncratic noise levels (operationalized through the Level 3 / TCE ratio and OTC derivative concentration) are systematically more vulnerable to noise-triggered failure, and they are more susceptible to the same systematic noise shock that affects all institutions simultaneously. Regulators should therefore monitor idiosyncratic noise levels during tranquil periods and require reductions through disclosure and asset reclassification before the systematic component of noise rises.

The dynamic macroprudential activation rule proposed in Section 7 operationalizes these theoretical insights as a concrete policy instrument. The η -gap $\Gamma_{i,t}$ and the systemic noise index Ξ_t provide regulators with a continuous, forward-looking indicator of market mechanism failure that complements the existing toolkit of leverage-based, liquidity-based, and systemic-risk-based triggers. The tiered intervention framework conditioned on $\Gamma_{i,t}$ provides the regulatory authority for preemptive action during the approach to the noise-dominated threshold, rather than the post-hoc crisis management that characterized the regulatory response to Bear Stearns in 2008.

Future research should pursue several extensions of the present framework. First, a formal estimation of the β coefficients in the noise proxy (22) from historical panel data on financial institution opacity and market signals would provide empirical validation of the calibration framework and enable rigorous out-of-sample testing. Second, the network externality extension discussed in Section 8.4, propagating idiosyncratic noise shocks through the financial network, would provide a richer theory of systemic noise amplification. Third, the feedback between η_t and the latent state variables V_t and λ_t should be modeled explicitly in a full equilibrium framework, allowing the welfare consequences of alternative regulatory intervention rules to be evaluated. Fourth, the international dimension of noise contagion (the channels through which elevated η_t^M in one jurisdiction transmits to others through global funding markets) is of immediate relevance to the design of international macroprudential coordination mechanisms under Basel III.

The 2008 financial crisis was not merely a failure of capital adequacy or of liquidity management, though it was both of those things. It was, at a deeper level, a failure of the informational infrastructure of markets. When the mechanism through which price signals convey information about fundamental values breaks down, even well-capitalized, solvent institutions can be swept into liquidation by the force of market panic. The framework developed in this paper and its predecessors provides a rigorous theoretical foundation for understanding this phenomenon and for designing the regulatory architecture required to prevent its recurrence.

References

1. Hanson, D. J. (2025). Evolving private information: A stochastic extension of informational asymmetry models. Working Paper, SSRN.
2. Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.
3. Leland, H. E. and Pyle, D. H. (1977). Informational asymmetries, financial structure, and financial intermediation. *Journal of Finance*, 32(2):371–387.
4. Myers, S. C. and Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2):187–221.
5. Agrawal, V. K. and Rao, R. K. S. (2015). *Corporate Policies in a World with Information Asymmetry*. World Scientific Publishing Co. Pte. Ltd., Singapore.
6. Hanson, D. J. (2026). Noise-dominated markets: Liquidity risk without insolvency. Working Paper, SSRN.
7. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417.
8. Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.
9. Daley, B. and Green, B. (2012). Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504.
10. Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.
11. Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates.
12. Duffie, D. and Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Re-view of Financial Studies*, 12(4):687–720.
13. Longstaff, F. A. and Schwartz, E. S. (1995). A simple approach to valuing risky fixed and floating rate debt. *Journal of Finance*, 50(3):789–819.
14. Collin-Dufresne, P., Goldstein, R. S., and Martin, J. S. (2001). The determinants of credit spread changes. *Journal of Finance*, 56(6):2177–2207.
15. Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3(1-2), 125-144.
16. Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101.
17. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384.
18. Ang, A. and Bekaert, G. (2002). Regime switches in interest rates. *Journal of Business and Economic Statistics*, 20(2):163–182.
19. Brendle, S. and Carmona, R. (2004). Hedging in incomplete markets. *SIAM Journal on Control and Optimization*, 43(6):2063–2083.
20. Xiong, W. (2001). Convergence trading with wealth effects: An amplification mechanism in financial markets. *Journal of Financial Economics*, 62(2):247–292.
21. Borio, C. (2003). Towards a macroprudential framework for financial supervision and regulation? *CESifo Economic Studies*, 49(2):181–215.
22. Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.
23. Hanson, S. G., Kashyap, A. K., and Stein, J. C. (2011). A macroprudential approach to financial regulation. *Journal of Economic Perspectives*, 25(1):3–28.
24. Tobias, A., & Brunnermeier, M. K. (2016). CoVaR. *The American Economic Review*, 106(7), 1705.
25. Gorton, G. B., & Metrick, A. (2012). Securitized banking and the run on repo.
26. Kashyap, A. K., Stein, J. C., & Hanson, S. (2010). An analysis of the impact of 'substantially heightened' capital requirements on large financial institutions. *Booth School of Business, University of Chicago, mimeo*,

27. Banerjee, R. N. and Mio, H. (2018). The impact of liquidity regulation on banks. *Journal of Financial Intermediation*, 35:30–44.
28. Bagehot, W. (1873). *Lombard Street: A Description of the Money Market*. Henry S. King & Co., London.
29. Rochet, J.-C. and Vives, X. (2004). Coordination failures and the lender of last resort: Was Bagehot right after all? *Journal of the European Economic Association*, 2(6):1116–1147.
30. Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–408.
31. Liu, J., & West, M. (2001). Combined parameter and state estimation in simulation-based filtering. In *Sequential Monte Carlo methods in practice* (pp. 197-223). New York, NY: Springer New York.
32. Schwartz, A. (2008). Interview on Squawk Box. CNBC, March 12.
33. Financial Crisis Inquiry Commission. (2011). *The Financial Crisis Inquiry Report, Authorized Edition: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the Un.* PublicAffairs.
34. Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–3153.
35. Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial. *Networks*, 105.
36. Cohan, W. D. (2009). *House of Cards: A Tale of Hubris and Wretched Excess on Wall Street*. Doubleday.
37. Cox, C. (2008a). Statement by SEC Chairman Christopher Cox Regarding the Bear Stearns Companies Inc. Securities and Exchange Commission Press Release.
38. Cox, C. (2008b). Letter to the Chairman of the Basel Committee on Banking Supervision. U.S. Securities and Exchange Commission.
39. Federal Reserve Bank of New York (2008). Statement on financing arrangement for the Bear Stearns Companies Inc. Press Release.
40. Federal Reserve Bank of New York (2018a). Maiden lane transactions. newyork-fed.org/markets/maidenlane.
41. Federal Reserve Bank of New York (2018b). New York Fed announces full repayment of Maiden Lane LLC loan. Press Release.
42. Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.
43. U.S. Securities and Exchange Commission, Office of Inspector General (2008). SEC’s Oversight of Bear Stearns and Related Entities. Technical Report 446-A.