# Narratives in Crisis and Optimal Manipulation Level for Public Media

Yuecheng He*
Punuo Zhang†

*School of Economics, Fudan University*

February 26, 2026

**Abstract**

This paper combines a global games model with textual data to analyze the role of narratives in crisis. We introduce an endogenous skepticism mechanism into a global games model to capture optimal manipulation of public signals. The model is calibrated to narratives and market returns in the 2015 Chinese stock market crisis. We use Qwen Local LLM to extract sentiments in news and *Guba* comments from all A-share stocks in 2015. We find that: First, narratives can predict the timing of a crisis. Second, there exists a coordination effect of public narratives in normal times, which amplifies their impact on market returns. Third, there exists an optimal level of manipulation for public narratives. Market performance first increases then decreases with manipulation level, exhibiting a non-monotonic relationship. Based on counterfactual analysis, we suggest public media should release more optimistic signals during a crisis.

*Email:yche24@m.fudan.edu.cn
†Email:pnzhang23@m.fudan.edu.cn

# 1 Introduction

Financial crises are rarely driven solely by fundamentals. Instead, they are often fueled by narratives disseminated through news media and social platforms. As Shiller (2017) says, "The human brain has always been highly tuned toward narratives, whether factual or not, to justify ongoing actions, even such basic actions as spending and investing." However, due to the lack of formal models to incorporate narratives into equilibrium analysis, and the difficulty in measuring narratives empirically, the role of narratives in crisis remains underexplored.

This paper develops a novel global games model with possible public signal manipulation to analyze the roles of public and private narratives in crisis, and uses LLM tools to measure narratives to calibrate the model. We show that narratives have a strong coordination effect on investors' behaviors. Also, when a crisis come, public narratives can be manipulated upwards to save the market. Such a manipulation has an optimal level due to a trade-off between level of signal and its credibility. The model is calibrated to data in the 2015 Chinese stock crisis, which perfectly predict the timing and magnitude of the crisis.

We develop a global games model to capture two key features of a crisis, multiple equilibria and strategic complementarity, and to describe how narratives can drive the market from a good equilibrium to a bad equilibrium. The model considers public narratives from news media that are observable to all investors and private narratives from individual investors' sentiments expressed on social platforms as public signals and private signals, respectively. It theoretically shows that narratives can affect market returns through their reflection of the true fundamental and coordination effects due to the information structure. We borrow the publicity multiplier from Morris and Shin (2003) to show that there exists a coordination effect of public narratives, which amplifies their impact on market returns.

In addition, we introduce a novel mechanism into the global games model which allows investors to endogenously decide whether to trust public narratives or not. We call it a hypothesis testing mechanism, supported by an axiomatic framework developed by Ortoleva (2012). The idea comes from an interesting observation in our textual data: private narratives turned more pessimistic after the market crash, but this did not happen to public narratives from news. One may interpret this pattern as public media trying to stabilize the market by manipulating public narratives to be more positive, and thus, biased upwards (relative to the true fundamental). Investors are not stupid and will re-evaluate the credibility of public signals. The endogenous skepticism mechanism contributes to the global games literature by allowing public signals to be potentially biased. It implies that there exists an optimal level of public narrative manipulation such that the aggregate holding ratio first increases then decreases with manipulation level, i.e., a non-monotonic relationship. Beyond the optimal point, the "Tacitus Trap" is triggered, where manipulated authoritative narratives that try to stabilize the market turn out to be counterproductive, as investors become more skeptical and tend to run.

To empirically test the model, I use the 2015 Chinese stock market crisis as an example. Figure 1 shows the Shanghai Stock Exchange Composite Index (SSE) during 2015. The market experienced a dramatic rise from around 3,000 points at the beginning of 2015 to a peak of 5,166 points on June 12, followed by a sharp crash back to below 3,000 points at


**Figure 1:** SSE Index in 2015. The dramatic rise started from early March, peaked on June 12, then crashed sharply till July 8, and finally experienced another drop in late August.

the end of August. The stock market in 2015 not only produces a remarkable boom and bust cycle, but also provides rich sources for studying narratives. During the year, SSE was driven up by a series of positive narratives such as ”Reform Bull” and ”New Economy”, and then crashed in the negative narratives like ”Malicious Shorting” and ”Deleveraging”. Current literature studies the 2015 stock crash in various perspectives, such as leverage-induced fire sales as major causes (Bian et al., 2025), wealth redistribution (An et al., 2022), and government interventions (Brunnermeier et al., 2022). However, the specific role of narratives in driving the market from boom to bust remains underexplored.

We collect news articles and *Guba* investor comments from all A-share stocks in 2015 as textual data, and use Qwen API to extract sentiments from the texts. Traditional textual analysis often struggles with context, irony, or the between-the-lines meaning of financial disclosures (Loughran and McDonald, 2011). Recent studies suggest that LLMs significantly outperform dictionary-based methods in textual analysis (e.g., Fang et al. (2025) in decoding industrial policies, and Lopez-Lira and Tang (2025) in predicting market returns). We then calibrate the model to the data and perform simulations to quantify the effects of narratives on market returns. The simulation results confirm the three main implications of the model: i) the equilibrium threshold, holding ratio, and skeptic ratio vary throughout 2015 in response to changing narratives, ii) the coordination effects of public narratives, measured by the publicity multiplier, fluctuate significantly. Days with the highest publicity multipliers coincide with major news releases such as the statement of ”Reform Bull” by the CSRC Chairman on March 10, 2015, and iii) a counterfactual analysis shows that the aggregate holding ratio first increases then decreases with manipulation level, confirming the non-monotonic relationship and the existence of an optimal manipulation level.

There are three main results of our model, later confirmed by calibration and simulation. First, narratives do affect market returns and can predict the timing of a crisis. It can be shown that signals people receive from news and social media comments influence their investment decisions. The pessimistic narratives emerged before the market crash. Second, in normal times, there exists a strong coordination effect of public narratives, which amplifies their impact on market returns. This result is the same as traditional global games models (Morris and Shin, 2003), where public signals not only reflect the true fundamental, but also trigger strategic coordination behaviors among investors. Third, and perhaps most interestingly, there exists an optimal level of manipulation that maximizes market performance. The aggregate holding ratio first increases and then decreases with the manipulation level, exhibiting a non-monotonic relationship. Beyond the optimal point. There is a trade-off between signal level and credibility. When the manipulation is small, investors think that bias between public and private signals is acceptable. The manipulation than pull up the posterior average of public and private signals. However, as the manipulation becomes too large, investors may not believe in that signal since it is too far away from their private signals. Due to the loss of the public signal's coordination power, investors are more likely to run. This is a novel mechanism we introduce into the global games model by letting public signals be potentially biased, and investors endogenously decide whether to believe the public signal or not through a criterion similar to hypothesis testing.

We use the calibrated parameter to analyze different counterfactual manipulation level. To our surprise, we find the realized manipulation level is lower than the optimal level, suggesting that public media can release more optimistic signals during a market crash. However, we admit that from an ex ante point of view, policymakers do not know the exact optimal manipulation level in the future, so such narrative support should be reluctant.

**Literature Review.** This paper contributes to three strands of literature. First, it connects to a large literature on financial crisis, particularly regarding the roles of investor beliefs and coordination failures. Classical theories of crises often emphasize the existence of multiple equilibria driven by self-fulfilling prophecies, where the shift from a good equilibrium to a bad one is triggered by sunspots (e.g., Diamond and Dybvig, 1983; Obstfeld, 1996; Gennotte and Leland, 1990). While these models capture the fragile nature of financial markets, they typically lack a specific mechanism to explain what triggers the shift in beliefs. More recent literature has turned to behavioral explanations, emphasizing how investors form beliefs based on extrapolation or diagnostic expectations (Barberis et al., 2018; Bordalo et al., 2018; Greenwood and Shleifer, 2014). In the context of the 2015 Chinese stock market crash, Bian et al. (2025) highlight the role of leverage-induced fire sales, yet the question of how narratives interact with these mechanisms remains underexplored.

Our paper complements this literature by offering a hypothesis testing mechanism that acts as an equilibrium selection device. Instead of relying on unexplained sunspots, we show that the interaction between public narratives and investors' trust can uniquely determine the timing of a crisis. This echoes the work on the social value of information (Angeletos and Pavan, 2007; Bond and Goldstein, 2015), but differs by focusing on the breakdown of trust in the information source itself.

Second, this paper relates to the extensive literature on global games models (Carlsson and van Damme (1993); Morris and Shin (2003); Frankel et al. (2003); He and Xiong (2012, e.g.,)). The global games framework demonstrates that by introducing a small amount of noise into signals, the multiple equilibria in coordination games can be pinned down, yielding a unique outcome determined by fundamental thresholds. Morris and Shin (2003) explicitly discuss the publicity multiplier effect, showing how public signals are more powerful than private ones in coordinating actions because they serve as a focal point for higher-order beliefs. This framework has been widely applied to bank runs (Goldstein and Pauzner, 2005), debt runs (He and Xiong, 2012), and currency crises (Angeletos et al., 2007).

To the best of our knowledge, the most relevant work to ours is Edmond (2013), who endogenizes the public signal by allowing a regime to manipulate its level or precision to withstand attacks. Our model is different from his in two means. First, his model assumes agents are standard Bayesians who account for the manipulation incentives but still update beliefs within the model's structure. Our paper departs from this by introducing endogenous skepticism through a "hypothesis testing" mechanism. Second, unlike Edmond (2013), where manipulation generally helps the regime, we show that excessive manipulation can trigger a "Tacitus Trap" where investors reject the public signal entirely. This aligns with recent efforts to model competing narratives (Eliaz and Spiegler, 2020) but frames it within a global games setting with potentially biased public signals.

Third, this paper contributes to the field of narrative economics and textual analysis in finance. Shiller (2017) famously argues that viral narratives can drive economic fluctuations independent of fundamentals. Empirically, early work by Antweiler and Frank (2004) and the seminal paper by Tetlock (2007) established the link between media content and market behavior. Tetlock (2007) uses a bag-of-words approach with the General Inquirer dictionary to show that high media pessimism predicts downward pressure on market prices. More recently, Goetzmann et al. (2024) use survey data to document crash narratives, finding that media reporting on rare disasters significantly heightens investors' perceived probability of a crash. Other works have linked news implied volatility to disaster concerns (Manela and Moreira, 2017) or utilized topic modeling to extract systematic risk factors (Bybee et al., 2023).

Our contribution to this literature is methodological. While most existing studies treat narratives as descriptive variables or regression predictors (e.g., García, 2013; Loughran and McDonald, 2011), we perform a model-based textual analysis. We map the sentiments extracted from news and social media directly into the parameters of a global games model as public signals and private signals, respectively. This allows for a quantitative evaluation of the coordination effects and the impact of narrative divergence, moving beyond reduced-form correlations.

Section 2 of this paper presents the model and derives the optimal manipulation result. Section 3 describes the textual data and sentiment extraction using LLM. Section **??** presents the calibration and simulation results. And finally, section 5 concludes the paper.

# 2 Model

## 2.1 Model Setup

Consider a model with a continuum of players, indexed by $i$. There is one type of asset in the market. Each player faces a binary choice: either to *Hold* or *Run*. Let $\theta$ denote the fundamental state of the market. In the context of the 2015 Chinese stock market crisis, $\theta$ represents the underlying market liquidity or the true economic health of the firms. Fur-

**Table 1:** Equilibrium Regions in the Complete Information Model.

| Fundamental | Region | Equilibrium Outcome |
|---|---|---|
| $\theta < 0$ | Lower Dominance Region | Unique (Run) |
| $0 \leq \theta \leq 1$ | Intermediate Region | Multiple (Run or Hold) |
| $\theta > 1$ | Upper Dominance Region | Unique (Hold) |

thermore, let $l$ denote the proportion of players who choose to *Hold*.

Each investor has a linear ex post payoff. The payoff to choosing *Run* is 0, while the payoff to choosing *Hold* is $\theta + l - 1$. This captures the strategic complementarity inherent in the market in the sense that more other investors choosing to hold ($l \uparrow$) increases one's payoff, and thus induces his choice to hold. On the other hand, if one believes other investors choose to run, it is the best choice for himself to run as well.

**Crisis Threshold.** Let $a(\theta) = 1 - l(\theta)$ denote the proportion of investors who choose to *Run*. We assume there exists an exogenous crisis threshold $a^* \in (0, 1)$ such that a crisis occurs if and only if $a(\theta) > a^*$. The threshold $a^*$ can be interpreted as the market's resilience to selling pressure. A crisis happens when the proportion of investors running exceeds this critical level. As we will show, equilibrium analysis pins down a fundamental threshold $\theta^*$ such that $a(\theta) > a^*$ if and only if $\theta < \theta^*$.

Suppose the true fundamental is common knowledge for now. Every investor knows $\theta$, as well as knowing the fact that other people know $\theta$, and so on. The payoff structure above gives the sufficient and necessary condition for any investor to choose to hold. For player $i$, he will choose to hold if and only if

$$\theta + l - 1 \geq 0, \text{ or } l \geq 1 - \theta. \tag{1}$$

Then the Nash equilibria are as follows. First, when $\theta > 1$, even if no one chooses to hold ($l = 0$), the inequality (1) holds. Therefore, all investors who observe a sufficiently high fundamental will hold. Call $\theta > 1$ the upper dominance region. Second, the lower dominance region of this game is defined by $\theta < 0$ because even if everyone chooses to hold ($l = 1$), the condition to hold in (1) still cannot hold. Therefore, in the lower dominance region, the best action for an investor is always to run.

In the intermediate region, however, there are multiple equilibria. When $\theta \in [0, 1]$, the best response of an investor is conditional on the actions of other investors. On the one hand, if an investor believes that everyone else will run, the payoff from holding is $\theta - 1 < 0$. Thus, he will also run. Everyone running ($l = 0$) is a self-fulfilling equilibrium. On the other, if he believes that everyone else will hold, the payoff from holding is $\theta > 0$. Thus, he will also hold. Everyone holding ($l = 1$) is also a self-fulfilling equilibrium. Table 1 concludes the equilibria under all three regions.

The simple model captures two main features of crisis: multiple equilibria and strategic complementarity. When prices are higher than the fundamental value in a stock market, it is possible that everyone is riding the bubble, believing that prices can go even higher. This is a bubble equilibrium. However, it is also possible that somehow some investors decide to run, selling shares they hold, leading to a crash. That is when a crisis happens. Traditional theories help explain these two features of crisis, but they cannot answer what drives the outcome from the good equilibrium to the bad equilibrium. Global games models are used to pin down multiple equilibria and specify forces that determine which equilibrium will happen. By incorporating narratives from news and social media into global games models, we believe our model can further identify the specific time point when the market goes from a bubble into a crisis.

### 2.1.1 Information Structure

The core of global games models (e.g., Morris and Shin (2003); Angeletos et al. (2007); Edmond (2013)) is that even a very little noise in signals that investors observe can break common knowledge and create huge strategic uncertainty, which helps to pin down a unique equilibrium outcome. In our model, the fundamental of the market cannot be observed directly by investors. They try to form beliefs about the fundamental based on what they see in the news and social media. All those sources of information can be categorized into public and private signals. Public signals are observable to all investors, while private signals are only observable to the investor who receives them.

Formally, assume the realization of the fundamental, $\theta$, is not directly observable. Instead, players observe a common public signal $y$ and heterogeneous private signals $x_i$, which are independent. Each player $i$ observes a private signal $x_i$ given by

$$x_i = \theta + \epsilon_i, \tag{2}$$

where $\epsilon_i \sim i.i.d. N(0, \sigma^2)$ is an idiosyncratic noise term independent across players. Later, this private signal $x_i$ is proxied by individual investor sentiment derived from comments on the *East Money (Guba) forum*.

Our biggest departure from current global games literature is that, we assume public signals can be manipulated by the public media. Suppose the public media introduces a bias $\Delta \geq 0$ into the public signal, such that

$$y = \theta + \Delta + \epsilon^y, \tag{3}$$

where $\epsilon^y \sim N(0, \tau^2)$. Notice that $\Delta$ is unobservable, but its potential existence is common knowledge. As a result, $y \sim N(\theta + \Delta, \tau^2)$. In the empirical part, $y$ corresponds to narratives, proxied by news sentiment from major media outlets.

Based on these signals, a naive investor's posterior belief regarding the fundamental $\theta$ would be a weighted average of the public and private signals, weighted by their precisions,

$$\bar{\theta}_i = \frac{\frac{1}{\tau^2} y + \frac{1}{\sigma^2} x_i}{\frac{1}{\tau^2} + \frac{1}{\sigma^2}} = \frac{\sigma^2 y + \tau^2 x_i}{\sigma^2 + \tau^2}. \tag{4}$$

However, investors are not naive. When the public signal appears to be too far from what they privately observe, they may suspect that the public signal is manipulated and choose not to trust it. This is the key departure of our model from the standard global games literature.

### 2.1.2 Endogenous Skepticism

To describe how people react to potentially manipulated public signals, it may come to one's mind naturally that investors refuse to believe public signals when they think the news is absurd, based on their own private signals. This, especially for economists, is very similar to hypothesis testing, in the sense that an investor has a null hypothesis that the public signal is unbiased, namely $H_0 : \Delta = 0$, against the alternative $H_A : \Delta > 0$. Notice that this is a one-sided test, because one only suspects that the public signal is manipulated upwards instead of downwards, especially in times of crisis.

Upon observing both the public signal $y$ and a private signal $x_i$, investor $i$ tests the null hypothesis that the public signal is unmanipulated ($H_0 : \Delta = 0$) against the alternative ($H_A : \Delta > 0$). The test statistic is the difference between the signals, $d_i = y - x_i$. Under $H_0$, $d_i \sim N(0, \sigma^2 + \tau^2)$. Given a significance level $\alpha$, the investor rejects the null hypothesis if

$$y - x_i > z_\alpha \sqrt{\sigma^2 + \tau^2}, \tag{5}$$

where $z_\alpha = \Phi^{-1}(1 - \alpha)$. Notice that $\alpha$ is an important parameter in the hypothesis testing mechanism, representing the investor's awareness of potential manipulation.

This mechanism of non-Bayesian updating is theoretically grounded in the literature on behavioral distortions and structural belief updates. Ortoleva (2012) proposes a "Hypothesis Testing model" to address the limitations of standard Bayes' rule in reacting to unexpected news. In his setup, an agent maintains a prior over priors and follows Bayes' rule until receiving an unexpected signal, which is defined as an event whose probability falls below a certain threshold. Upon such realization, the initial prior fails the test, and the agent experiences a paradigm shift, rejecting her current prior to select a new one that maximizes the updated likelihood.

We apply this axiomatic framework to a financial and macroeconomic setting to capture aggregate market behavior and systemic narrative collapse, with a slightly different interpretation of the threshold. In Ortoleva (2012), the threshold is a subjective behavioral parameter measuring the degree of dynamic inconsistency, which he sets to zero to analyze the limiting scenario. Our model, however, uses $\alpha$ to capture not only micro-level behavioral biases but also macro-level aggregate intolerance to biases in public signals. As aforementioned, a larger $\alpha$ means that investors are more sensitive to potential manipulations, which means they are more cautious. Unlike economists who usually set $\alpha = 0.05$, investors in our model adopt a much higher threshold (e.g., around 20% in our simulations) to reflect their systemic intolerance toward potential information manipulation during crisis. (Supported by Rabin and Schrag (1999); Gentzkow and Shapiro (2006)).

This mechanism endogenously partitions the population into two types:

1. *Skeptics* (Type S): the investors with $x_i < \bar{x}(y)$, where $\bar{x}(y) \equiv y - z_\alpha \sqrt{\sigma^2 + \tau^2}$. They reject $H_0$, deem the public signal unreliable, and discard $y$, relying solely on $x_i$.
2. *Believers* (Type B): the investors with $x_i \geq \bar{x}(y)$. They do not reject $H_0$ and update beliefs using both $x_i$ and $y$, as in the standard Bayesian updating.

## 2.2 Equilibrium

### 2.2.1 Posterior Beliefs

The two types of investors form different posterior beliefs about the fundamental $\theta$. A Skeptic discards the public signal $y$ and relies solely on her private signal $x_i$. Her posterior belief is

$$\theta | x_i \sim N(x_i, \sigma^2) \quad \text{for Skeptics.} \tag{6}$$

A Believer, on the other hand, incorporates both signals. Her posterior belief follows the standard Bayesian updating formula:

$$\theta | x_i, y \sim N\left(\frac{\sigma^2 y + \tau^2 x_i}{\sigma^2 + \tau^2}, \frac{\sigma^2 \tau^2}{\sigma^2 + \tau^2}\right) \quad \text{for Believers.} \tag{7}$$

Denote the Believer's posterior mean by $\overline{\theta}_i = \frac{\sigma^2 y + \tau^2 x_i}{\sigma^2 + \tau^2}$ and the posterior variance by $\tilde{\sigma}^2 = \frac{\sigma^2 \tau^2}{\sigma^2 + \tau^2}$. Notice that Believers have more precise posteriors than Skeptics ($\tilde{\sigma}^2 < \sigma^2$), reflecting the additional information from the public signal.

For a given fundamental $\theta$, denote the fraction of Skeptics in the population by $\eta(\theta, y) \equiv \Phi\left(\frac{\bar{x}(y) - \theta}{\sigma}\right)$, where $\Phi(\cdot)$ is the standard normal CDF. As $y$ increases (either due to higher $\theta$ or higher manipulation $\Delta$), the skepticism cutoff $\bar{x}(y)$ rises, and more investors become Skeptics.

### 2.2.2 Threshold Strategy

Following the global games literature, we consider monotone switching strategies. The presence of two investor types—Skeptics and Believers—requires a generalized approach. We characterize the equilibrium by two cutoffs $(\kappa_S, \kappa_B)$: a Skeptic chooses to *Hold* if and only if $x_i > \kappa_S$, while a Believer chooses to *Hold* if and only if $x_i > \kappa_B$. Since the two types have different posterior beliefs, it is convenient to express their strategies in terms of private signals rather than posterior beliefs.

The aggregate proportion of investors choosing to *Hold*, denoted by $L(\theta, y)$, combines contributions from both types. Since private signals $x_i$ are distributed as $N(\theta, \sigma^2)$, the mass of each type holding can be computed directly. Skeptics hold if $\kappa_S < x_i < \bar{x}(y)$; Believers hold if $x_i > \max(\kappa_B, \bar{x}(y))$. Summing these contributions:

$$L(\theta, y) = \underbrace{\left[\Phi\left(\frac{\bar{x}(y) - \theta}{\sigma}\right) - \Phi\left(\frac{\kappa_S - \theta}{\sigma}\right)\right]^+}_{\text{Skeptics who hold}} + \underbrace{1 - \Phi\left(\frac{\max(\kappa_B, \bar{x}(y)) - \theta}{\sigma}\right)}_{\text{Believers who hold}}, \tag{8}$$

where $[z]^+ = \max(z, 0)$. This tractable expression shows that $L(\theta, y)$ is monotonically increasing in $\theta$, preserving the strategic complementarity necessary for equilibrium uniqueness.

Given the payoff function $\theta + l - 1$, the equilibrium thresholds $(\kappa_S, \kappa_B)$ are determined by indifference conditions for marginal investors. The marginal Skeptic (with $x_i = \kappa_S$) has

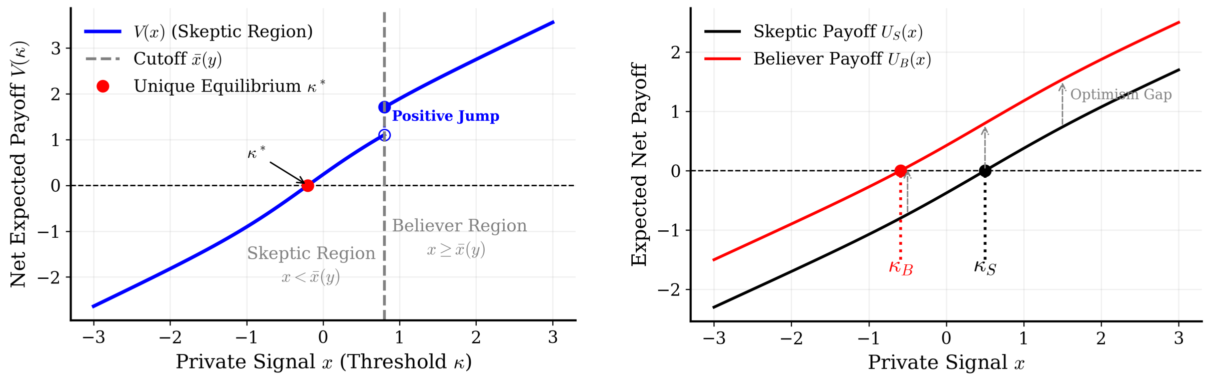


**Panel A:** Uniqueness of Equilibrium

**Panel B:** $\kappa_S > \kappa_B$

**Figure 2:** Sketches for Proposition 2. Panel A demonstrates the uniqueness of the equilibrium and Panel B illustrates the ordering of the thresholds.

posterior $\theta \sim N(\kappa_S, \sigma^2)$. Her expected payoff from holding equals zero when:

$$\mathbb{E}[\theta + L(\theta, y) - 1 | x_i = \kappa_S] = \kappa_S + \mathbb{E}[L(\theta, y)|\kappa_S] - 1 = 0. \tag{9}$$

The marginal Believer (with $x_i = \kappa_B$) has posterior $\theta \sim N(\bar{\theta}_B, \tilde{\sigma}^2)$, where $\bar{\theta}_B = \rho\kappa_B + (1 - \rho)y$ and $\rho = \frac{\tau^2}{\sigma^2+\tau^2}$. His indifference condition is:

$$\mathbb{E}[\theta + L(\theta, y) - 1 | x_i = \kappa_B, y] = \bar{\theta}_B + \mathbb{E}[L(\theta, y)|\kappa_B, y] - 1 = 0. \tag{10}$$

**Proposition 1 (Existence and Uniqueness of Equilibrium).** For any public signal $y$ and significance level $\alpha$, there exists a unique equilibrium characterized by thresholds $(\kappa_S^*, \kappa_B^*)$ with $\kappa_S^* > \kappa_B^*$.

The equilibrium can be equivalently characterized by a cutoff $\kappa^*$ such that all investors with $x_i > \kappa^*$ hold. The cutoff satisfies $V(\kappa^*) = 0$, where

$$V(x) = \begin{cases} x + \mathbb{E}[L|x] - 1, & \text{if } x < \bar{x}(y) \quad \text{(Skeptic)} \\ \bar{\theta}(x) + \mathbb{E}[L|x, y] - 1, & \text{if } x \geq \bar{x}(y) \quad \text{(Believer)}, \end{cases} \tag{11}$$

with $\bar{\theta}(x) = \rho x + (1 - \rho)y$.

The proof is in Appendix A. A graphical illustration of the proposition is in Figure 2.

The intuition is as follows. In standard global games, the value function $V(x)$ is continuous and monotonically increasing, so there is a unique cutoff. In our model, $V(x)$ has a discontinuous upward jump at $\bar{x}(y)$—crossing from the Skeptic to the Believer region raises the posterior mean from $x$ to $\rho x + (1 - \rho)y > x$ (since $y > \bar{x}(y)$). Despite this jump, monotonicity within each region and the boundary behavior ($V \to -\infty$ as $x \to -\infty$, $V \to +\infty$ as $x \to +\infty$) guarantee a unique crossing.

### 2.2.3 Equilibrium Threshold

With the signal threshold $\kappa^*$ determined, we can now characterize when a crisis occurs. Recall that the proportion of investors choosing to run is $a(\theta) = 1 - L(\theta, y)$, and a crisis occurs when $a(\theta) > a^*$, or equivalently, $L(\theta, y) < 1 - a^*$.

Since $L(\theta, y)$ is monotonically increasing in $\theta$, there exists a unique fundamental threshold $\theta^*$ such that $L(\theta^*, y) = 1 - a^*$. For $\theta < \theta^*$, we have $L(\theta, y) < 1 - a^*$, so $a(\theta) > a^*$ and a crisis occurs. For $\theta > \theta^*$, we have $L(\theta, y) > 1 - a^*$, so $a(\theta) < a^*$ and no crisis occurs.

**Proposition 2 (Existence and Uniqueness of Crisis Threshold).** Given the unique signal threshold $\kappa^*$ from Proposition 1, there exists a unique fundamental threshold $\theta^* \in \mathbb{R}$ such that:

1. If $\theta < \theta^*$, then $a(\theta) > a^*$ and a crisis occurs.
2. If $\theta > \theta^*$, then $a(\theta) \leq a^*$ and no crisis occurs.

The threshold $\theta^*$ is implicitly defined by $L(\theta^*, y) = 1 - a^*$.

*Proof.* By Proposition 1, the equilibrium threshold $\kappa^*$ exists and is unique. Given $\kappa^*$, the aggregate holding function $L(\theta, y)$ is well-defined and strictly increasing in $\theta$, with $\lim_{\theta \to -\infty} L(\theta, y) = 0$ and $\lim_{\theta \to +\infty} L(\theta, y) = 1$. Since $1 - a^* \in (0, 1)$, by the intermediate value theorem, there exists a unique $\theta^*$ satisfying $L(\theta^*, y) = 1 - a^*$. The crisis characterization follows directly from the monotonicity of $L$ in $\theta$. □

#### 2.2.4 Comparative Statics

We now analyze how the crisis threshold $\theta^*$ responds to changes in model parameters. Recall that a lower $\theta^*$ implies greater market stability as crises only occur when fundamentals fall below this threshold.

**Effect of Significance Level $\alpha$.** The significance level governs how readily investors reject the public signal. Notice that $z_\alpha = \Phi^{-1}(1 - \alpha)$ is *decreasing* in $\alpha$: a higher $\alpha$ (more willing to reject $H_0$) means a lower critical value. This raises the skepticism cutoff $\bar{x}(y) = y - z_\alpha \sqrt{\sigma^2 + \tau^2}$, increasing the mass of Skeptics. Therefore,

$$\frac{\partial \theta^*}{\partial \alpha} > 0. \tag{12}$$

More skeptical investors (higher $\alpha$) lead to a higher crisis threshold and a more fragile market, because more investors ignore the public signal and cannot be coordinated.

**Effect of Crisis Threshold $a^*$.** A higher $a^*$ means the market tolerates more selling pressure before a crisis is triggered. Since $L(\theta^*, y) = 1 - a^*$, a higher $a^*$ requires a lower $L$, which corresponds to a lower $\theta^*$:

$$\frac{\partial \theta^*}{\partial a^*} < 0. \tag{13}$$

**Effect of Signal Precisions.** The effects of $\sigma^2$ and $\tau^2$ on $\theta^*$ involve multiple channels and are generally ambiguous. Consider an increase in private signal precision (decrease in $\sigma^2$). On one hand, Believers place less weight on the public signal $y$ in forming posteriors—the weight $\frac{\sigma^2}{\sigma^2+\tau^2}$ falls—weakening the coordination effect. Simultaneously, the skepticism cutoff $\bar{x}(y) = y - z_\alpha \sqrt{\sigma^2 + \tau^2}$ rises, and the tighter distribution of private signals around

$\theta$ means more investors find themselves below this cutoff. Both forces expand the mass of Skeptics and reduce the power of public communication, tending to raise $\theta^*$ and destabilize the market. On the other hand, more precise private signals reduce strategic uncertainty: each investor has a better estimate of $\theta$, which can improve coordination through the private information channel. This force tends to lower $\theta^*$. The net effect depends on which channel dominates and is parameter-dependent. Similar ambiguity applies to changes in public signal precision $\tau^2$, where the weight, skepticism, and informativeness channels again pull in different directions.

## 2.3 Coordination Effects

In this section, we explore how public signals affect equilibrium beyond their pure informational content. The analysis follows Morris and Shin (2003) and focuses on Believers, who incorporate the public signal $y$ into their decisions. Since $\Delta$ is a constant bias that does not affect variances, the coordination effects derived below are structurally unchanged by manipulation.

Consider a Believer with posterior belief $\overline{\theta}_i = \frac{\sigma^2 y + \tau^2 x_i}{\sigma^2+\tau^2}$. In a symmetric equilibrium where all Believers use cutoff $\kappa_B$, the marginal Believer's indifference condition can be written as:

$$\overline{\theta} + \mathbb{E}[L(\theta, y)|\overline{\theta}] - 1 = 0, \tag{14}$$

where $\overline{\theta} = \frac{\sigma^2 y + \tau^2 x}{\sigma^2+\tau^2}$ is her posterior mean. Let $z \equiv \sqrt{\gamma}\,(\overline{\theta} - y) = \sqrt{\gamma}\left(\frac{\sigma^2 y + \tau^2 x}{\sigma^2+\tau^2} - y\right)$. Totally differentiating with respect to $y$ gives:

$$\frac{\sigma^2}{\sigma^2+\tau^2} + \frac{\tau^2}{\sigma^2+\tau^2}\frac{dx}{dy} - \phi(z)\sqrt{\gamma}\left(\frac{\sigma^2}{\sigma^2+\tau^2} + \frac{\tau^2}{\sigma^2+\tau^2}\frac{dx}{dy} - 1\right) = 0. \tag{15}$$

Solving for $\frac{dx}{dy}$,

$$\frac{dx}{dy} = -\frac{\frac{\sigma^2}{\tau^2} + \sqrt{\gamma}\phi(z)}{1 - \sqrt{\gamma}\phi(z)}. \tag{16}$$

This expression represents exactly the amount the private signal would have to change to compensate for a unit change in the public signal to keep the marginal investor indifferent between *Hold* and *Run*.

Next, we separate the pure information effects of the public signal in locating the fundamental $\theta$. That is, fixing the posterior belief $\overline{\theta}$, one asks how much her private signal must change to compensate for a given change in $y$. Totally differentiating the posterior belief equation $\overline{\theta} = \frac{\sigma^2 y + \tau^2 x}{\sigma^2+\tau^2}$ while treating $\overline{\theta}$ as a constant yields

$$\frac{dx}{dy} = -\frac{\sigma^2}{\tau^2}. \tag{17}$$

**Definition.** Define the *relative precision* as

$$\gamma \equiv \frac{\sigma^2}{\tau^4}\left(\frac{\sigma^2+\tau^2}{\sigma^2+2\tau^2}\right). \tag{18}$$

The *publicity multiplier* $\zeta$ is defined as the ratio of the total effects to the pure information effects:

$$\zeta = \frac{1 + \frac{\tau^2}{\sigma^2}\sqrt{\gamma}\phi(z)}{1 - \sqrt{\gamma}\phi(z)}. \tag{19}$$

This multiplier measures the scale of the public signal's strategic effects, or *coordination effects*, which influence players' strategic behavior in coordination (the decision to run or hold based on others' actions) rather than merely their beliefs about the fundamental $\theta$.

For the marginal investor whose posterior belief is $\overline{\theta}^* = \kappa$, the publicity multiplier is

$$\zeta^* = \frac{1 + \frac{\tau^2}{\sigma^2}\sqrt{\gamma}\phi(\Phi^{-1}(\overline{\theta}^*))}{1 - \sqrt{\gamma}\phi(\Phi^{-1}(\overline{\theta}^*))}. \tag{20}$$

Notice that the more precise or concentrated the public signal is versus the private signals (i.e., higher $\gamma$), the greater the publicity multiplier $\zeta^*$ will be. This suggests that during normal times, public narratives may have a disproportionately large impact on investor behavior compared to dispersed individual sentiments. However, as shown in the previous section, when the manipulation $\Delta$ becomes too large, investors turn into Skeptics and the coordination effect collapses.

## 2.4 Optimal Manipulation

A natural question arises: given that public signals can be manipulated, what is the optimal level of manipulation from the perspective of market stability? The conventional wisdom suggests that any manipulation is harmful, as it distorts information and undermines market efficiency. However, our model reveals a more nuanced picture. The relationship between manipulation and market stability is non-monotonic: a moderate level of manipulation can actually stabilize the market, while excessive manipulation backfires.

### 2.4.1 The Trade-off

When the public media increases the manipulation level $\Delta$, two opposing forces come into play.

**The Coordination Benefit.** A higher $\Delta$ raises the public signal $y = \theta + \Delta + \epsilon^y$. For Believers who trust the public signal, a higher $y$ shifts their posterior beliefs upward:

$$\overline{\theta}_i^B = \frac{\sigma^2 y + \tau^2 x_i}{\sigma^2 + \tau^2} = \frac{\sigma^2(\theta + \Delta + \epsilon^y) + \tau^2 x_i}{\sigma^2 + \tau^2}. \tag{21}$$

This makes Believers more optimistic about the fundamental, increasing their propensity to hold. Moreover, since public signals serve as coordination devices, a more optimistic public narrative can coordinate Believers toward the good equilibrium where everyone holds.

**The Credibility Cost.** On the other hand, a higher $\Delta$ also raises the skepticism cutoff $\bar{x}(y) = y - z_\alpha\sqrt{\sigma^2 + \tau^2}$. This means that for any given private signal $x_i$, the gap $y - x_i$

widens. More investors find the public signal implausibly optimistic relative to their private information, leading them to reject it. Formally, the mass of Skeptics increases:

$$\eta(\theta, y) = \Phi\left(\frac{\bar{x}(y) - \theta}{\sigma}\right) = \Phi\left(\frac{y - z_\alpha\sqrt{\sigma^2 + \tau^2} - \theta}{\sigma}\right). \tag{22}$$

Since $\partial y/\partial\Delta = 1$, we have $\partial\eta/\partial\Delta > 0$. These newly converted Skeptics discard the public signal entirely and rely solely on their private signals. They cannot be coordinated by any public narrative, no matter how optimistic.

### 2.4.2 Non-Monotonicity of Aggregate Holding

To state our main result on the non-monotonicity of manipulation, we introduce one additional condition on the baseline informativeness of the skepticism mechanism.

**Assumption 1 (Baseline Credibility).** At zero manipulation ($\Delta = 0$ and $\epsilon^y = 0$), the skepticism cutoff lies below the Believer equilibrium threshold:

$$\bar{x}(y)\big|_{\Delta=0} = \theta - c < \kappa_B^*(0), \tag{23}$$

where $c \equiv z_\alpha\sqrt{\sigma^2 + \tau^2}$.

This assumption says that when there is no manipulation, the unbiased public signal is credible to the marginal investor. Specifically, the marginal Believer (with private signal $x_i = \kappa_B^*$) does not find the public signal suspicious enough to reject. Equivalently, the skepticism tolerance band $c$ is wide enough that the marginal investor's observed discrepancy $y - x_i$ falls within the acceptance region of the hypothesis test. This is a natural baseline condition: if investors are already massively skeptical when the public signal is truthful, the concept of manipulation "backfiring" loses its meaningful starting point.

The assumption is satisfied whenever $\alpha$ is not too close to $\frac{1}{2}$ (investors are not too skeptical) and the signal noise $\sigma^2 + \tau^2$ is not too small (the tolerance band $c = z_\alpha\sqrt{\sigma^2 + \tau^2}$ is sufficiently wide). In our simulations with $\alpha \approx 0.20$, the condition holds comfortably for all parameter configurations considered.

**Proposition 3 (Non-Monotonicity of Manipulation).** Under the assumption $\tilde{\sigma}^2 \equiv \frac{\sigma^2\tau^2}{\sigma^2+\tau^2} > \frac{1}{2\pi}$ and Assumption 1, fix $\theta \in (0, 1)$ and condition on $\epsilon^y = 0$. The aggregate holding mass $\bar{L}(\theta; \Delta)$ satisfies:

(i) *(Initially increasing.)* There exists $\underline{\Delta} > 0$ such that $\bar{L}(\theta; \Delta) > \bar{L}(\theta; 0)$ for all $\Delta \in (0, \underline{\Delta})$.

(ii) *(Convergence and the Tacitus Trap.)*

  (a) For all $\theta \in (0, 1)$, $\lim_{\Delta\to\infty} \bar{L}(\theta; \Delta) = \bar{L}_S(\theta) \equiv 1 - \Phi\left(\frac{\kappa_S^\infty - \theta}{\sigma}\right)$, where $\kappa_S^\infty$ is the equilibrium threshold of the pure private-information game.

  (b) For $\theta > \kappa_S^\infty$, $\bar{L}_S(\theta) < \bar{L}(\theta; 0)$: excessive manipulation eventually drives the holding mass *below* the unmanipulated level. This is the *Tacitus Trap*.

(iii) *(Interior maximum.)* For all $\theta \in (0, 1)$, there exists $\Delta^* > 0$ such that $\bar{L}(\theta; \Delta^*) = \max_{\Delta\geq 0} \bar{L}(\theta; \Delta)$, and $\bar{L}(\theta; \Delta^*) > \max\{\bar{L}(\theta; 0),\ \bar{L}_S(\theta)\}$.

The proof is in Appendix B. See Figure 10 in our simulation for a graphical illustration.

When $\Delta$ is small, Assumption 1 ensures that essentially all investors trust the public signal and are Believers. A marginal increase in $\Delta$ raises the public signal $y$, which shifts Believers' posteriors upward and lowers the equilibrium threshold through the coordination effect, inducing strictly more investors to hold. This establishes part (i). As $\Delta \to \infty$, the public signal becomes unconvincingly optimistic and every investor rejects it, becoming a Skeptic who relies solely on private information. The equilibrium then converges to a pure private-information game. Since Skeptics lack the coordination benefit of public signals, their equilibrium threshold is strictly higher ($\kappa_S > \kappa_B$ by Proposition 1), and the resulting holding mass is strictly lower. This establishes part (ii). Part (iii) follows from parts (i) and (ii) together with the continuity of $\bar{L}$ in $\Delta$

#### 2.4.3 The Tacitus Trap Revisited

The non-monotonicity result provides a formal characterization of what we call the *Tacitus Trap*—a political science concept named after the Roman historian Tacitus, who observed that when authority loses the trust of the public, any attempt to assert control or disseminate information becomes futile, regardless of the content.

In our model, the Tacitus Trap manifests when manipulation becomes excessive. Beyond a certain point, attempts to stabilize the market through optimistic narratives are not only ineffective but actively harmful. The more the public media tries to reassure investors, the more investors become suspicious. This creates a vicious cycle: aggressive manipulation breeds skepticism, which undermines the coordination effect of public signals, which in turn may prompt even more aggressive manipulation.

The Tacitus Trap equilibrium is characterized by a high mass of Skeptics ($\eta \to 1$) and a collapsed publicity multiplier ($\zeta \to 1$). In this regime, the public signal loses its strategic coordination power and only retains its informational role for the few remaining Believers. The market becomes fragmented, with Skeptics and Believers forming beliefs based on different information sets, leading to greater uncertainty and potentially more severe crises.

#### 2.4.4 Policy Implications

Our analysis suggests several policy implications for market regulators and public media during financial crises.

First, there exists an optimal level of positive spin that can help coordinate investors and prevent panic runs. Complete transparency ($\Delta = 0$) is not necessarily optimal when coordination failures are the primary concern. A moderate amount of optimism in official communications can serve as a focal point for investor coordination.

Second, the effectiveness of manipulation depends on model parameters. When private signals are more precise (lower $\sigma^2$), investors rely more on their own information and the coordination benefit from public signals diminishes. When investors are more skeptical (higher $\alpha$), manipulation is more likely to backfire. Regulators should calibrate their communication strategies based on the prevailing market conditions and investor sentiment.

Third, once the Tacitus Trap is triggered, it is difficult to escape. The public media's credibility, once lost, cannot be easily restored. This suggests that regulators should be cautious about aggressive manipulation, especially during the early stages of a crisis when investor trust is still intact.

# 3 Text as Data

To incorporate narratives into the above model, we use text-as-data approaches to extract sentiment and certainty from news articles and social media comments on Chinese stock market in the whole year of 2015, during which the Chinese stock market experienced a significant crisis. News, we concentrate on authoritative media sources, represents public signals that are observable by all investors. On the other hand, comments from individual investors on social media platforms reflect private signals that vary across investors. We construct the 2015 Return–Narrative panel by integrating public news narratives, private forum narratives, and daily stock market data under a unified trading-day alignment rule. The raw data are first cleaned at the record level and then aggregated into market-day and stock-day signals.

## 3.1 Data Sources

For public signals, We collect 83,714 pieces of news from January 1, 2015 to December 31, 2015. The news are from 251 different sources, including major authoritative media outlets such as *Shanghai Securities News*, a leading state-run financial daily affiliated with Xinhua News Agency and serves as a designated authoritative channel for statutory information disclosure by listed companies, and famous non-authoritative media such as *21st Century Business Herald*, a well-known private financial newspaper in China. Table 2 lists the top 10 news sources by article counts. However, we only use the news from state-media and government-media as the proxy of our public signal in subsequent calibration, while we still use all news in this section for a whole picture visualization.

For private signals, it is hard to detect the noise of each individual investor's signal directly, as is defined by the "private" nature of these signals. However, it's neccesary to simplify that investors who receive a private signal and then express their opinions on social media platforms have some consistency in what they see and think. Assume perfect consistency, then comments on social media by individual investors can be treated as a proxy for private signals.

We collect 4,978,518 pieces of comments from a popular Chinese stock market social media platform, *East Money (Guba) forum*, during 2015. *East Money (Guba) forum* is one of the largest online investment communities in China, with millions of users discussing stock market trends, sharing investment strategies, and expressing their sentiments about various stocks and market conditions. *Guba* has a specific room for each listed company, which means we could map every comment to the targeted stock.

To avoid look-ahead bias, posts or news released after 15:00 or on non-trading days are mapped to the next trading day.

**Table 2:** Top 10 News Sources by Article Counts.

| Newspaper | Count | Newspaper | Count |
|---|---|---|---|
| Securities Times | 9,847 | Economic Information Daily | 3,855 |
| Shanghai Securities News | 9,060 | Beijing Business Today | 3,849 |
| National Business Daily | 8,287 | Yicai | 3,703 |
| 21st Century Business Herald | 8,075 | Financial Investment News | 3,506 |
| Securities Daily | 7,230 | Securities Times Online | 2,859 |

## 3.2 Sentiments and Convictions

There is much information embedded in the text data. Based on the model we established, the text data analysis aims to extract two key dimensions from the text—sentiment and conviction—to evaluate narratives in crisis numerically. Sentiment means how positive or negative the text is towards the stock market. Conviction distinguishes how strongly the text conveys its sentiment. For example, ”The market is probably doing well the next morning” and ”Definitely it's booming!” both express positive sentiments, but the latter shows a higher level of conviction and thus will weight more in the analysis. Sentiments in news will represent public signals $y$, while sentiments in social media comments will represent private signals $x$. As we've explained in the front context that the private signal is able to be mapped into specific stock, we refer it as $x_{i,t}$, where i represents stock and t represents time. However, public signal is more created to push the belief of fundamental but specific stock, we refer it as $y_t$.

Another purpose of this exercise is to provide information regarding precisions of signals, namely $\tau$ and $\sigma$. Since LLMs provides sentiment scores and conviction scores for each piece of text independently, we can then calculate the dispersion of sentiments weighted by convictions to obtain their variances and standard deviations.

Since we've posed the detailed Promopts in the Appendix, here we just discuss some key ideas for following process. First, we set the range of sentiment scores from -1 to 1, representing extremely negative sentiment to extremely positive sentiment. Second, we set the range of conviction scores from 0 to 1, where 0 indicates no conviction and 1 indicates extremely confident. Lastly, to make it comparable across news and comments, we keep the prompt as similar as possible, except for some necessary adjustments (e.g., changing objective ”good news” to subjective ”positive opinions” from news to comments).

Using certainty as the weight, we derived $x_{i,t}$ and $y_t$ as following:

$$x_{i,t} = \frac{\sum_{d \in \mathrm{D}_{i,t}^{forum}} s_d c_d}{\sum_{d \in \mathrm{D}_{i,t}^{forum}} c_d}.$$

$$y_t = \frac{\sum_{d \in \mathrm{D}_{t}^{news}} s_d c_d}{\sum_{d \in \mathrm{D}_{t}^{news}} c_d}.$$

Table 3 summarizes the statistics of news articles and social media comments used in this paper. The last rows of each panel present the weighted sentiment scores, calculated as the product of sentiment score and its conviction.

**Table 3:** Summary Statistics of News and Comments.

| **Panel A: News** | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Variable | Obs | Mean | Std | Min | Max | P25 | P50 | P75 |
| # /day (all days) | 365 | 229.35 | 142.21 | 0 | 454 | 93 | 280 | 350 |
| # /day (trd. days) | 244 | 320.98 | 59.11 | 105 | 454 | 280 | 324 | 365.5 |
| Words/article | 83 714 | 3 182.54 | 3 680.90 | 0 | 243 949 | 1 517 | 2 501 | 3 837 |
| Sentiment $s$ | 83 675 | 0.09 | 0.55 | -1.00 | 1.00 | 0.00 | 0.00 | 0.00 |
| Certainty $c$ | 83 675 | 0.87 | 0.22 | 0.00 | 1.00 | 0.50 | 1.00 | 1.00 |
| $s \times c$ | 83 675 | 0.07 | 0.49 | -1.00 | 1.00 | 0.00 | 0.00 | 0.00 |
| **Panel B: Comments** | | | | | | | | |
| Variable | Obs | Mean | Std | Min | Max | P25 | P50 | P75 |
| # /day (all days) | 365 | 13 639.78 | 8 402.47 | 684 | 32 999 | 4 223 | 14 295 | 20 966 |
| # /day (trd. days) | 244 | 18 770.70 | 5 005.33 | 8 118 | 32 999 | 14 259.8 | 18 826.5 | 22 779.5 |
| Words/comment | 4 978 518 | 67.93 | 261.16 | 0 | 23 309 | 10 | 17 | 29 |
| Sentiment $S$ | 4 978 518 | -0.04 | 0.65 | -1.00 | 1.00 | -0.50 | 0.00 | 0.20 |
| Conviction $C$ | 4 978 518 | 0.49 | 0.31 | 0.00 | 1.00 | 0.30 | 0.40 | 0.70 |
| $S \times C$ | 4 978 518 | -0.02 | 0.49 | -1.00 | 1.00 | -0.25 | 0.00 | 0.02 |

*P25/P50/P75 denote 25th/50th/75th percentiles. Daily counts shown for all calendar days and trading days; words are token counts per item; sentiments in* $[-1, 1]$ *and convictions in* $[0, 1]$.

### 3.2.1 Comments vs. News Sentiments

It is necessary to compare the signals extracted from news articles and social media comments. Figure 3 presents two panels showing the trends of sentiments in 2015. Both panels contain only trading days to eliminate the holiday effects (e.g., during the Spring Festival, comments were quite positive). It is obvious that sentiments from news articles (Panel A) are more stable and concentrated, while sentiments from Guba comments (Panel B) are more volatile and dispersed. This observation implies a high $\gamma$, i.e., high relative precision of public signals over private signals, in the model.

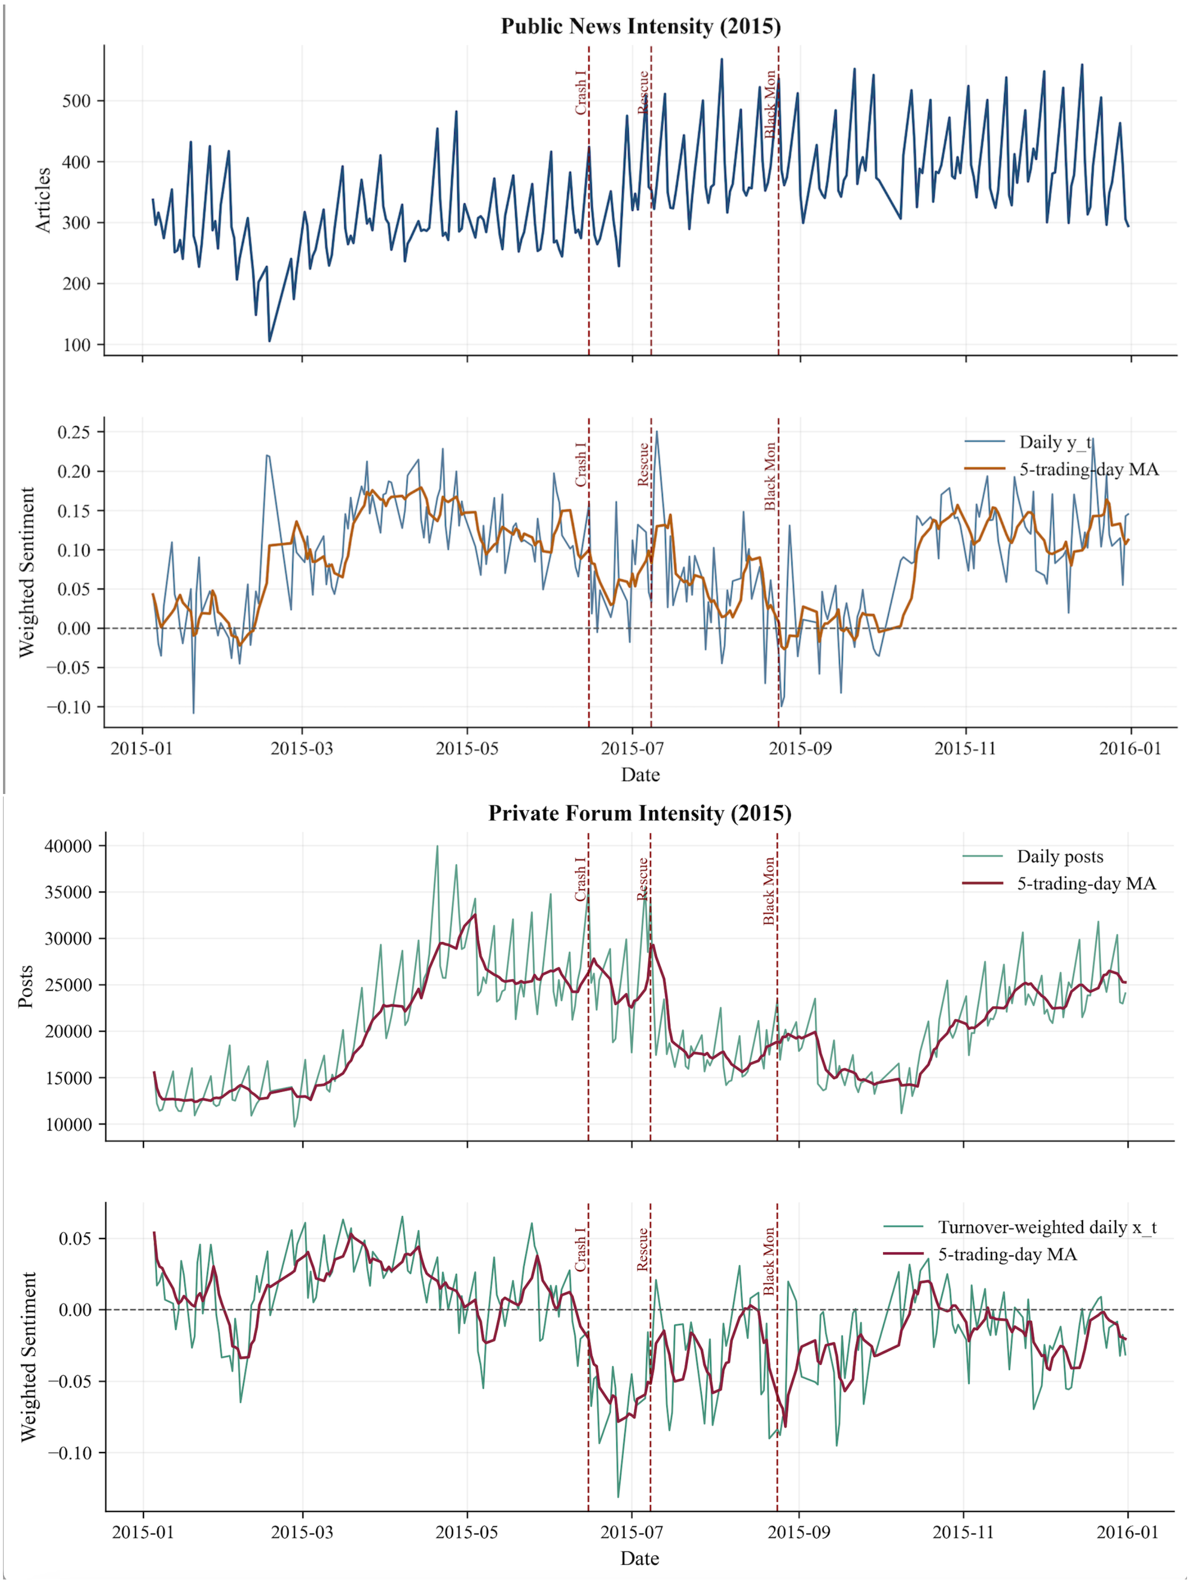


**Figure 3:** The intensity and sentiments of public signals (Panel A) and private signals (Panel B) during 2015.

In addition, focusing on the period around June 12, when the SSE Index peaked and then started to crash, one can see that news sentiments almost did not change, except for a small increase right after the crash (One may consider this as a manipulated counter-cyclical narrative intended to save the market from collapse). However, comments sentiments tell a different story. A few days before June 12, while the market was still booming, comments sentiments had already started to decline to the nearly lowest level up to that point. That is to say, individual investors had already sensed something wrong with the market through their private signals, even though the public narratives from news were still optimistic. In some sense, the pre-crash divergence between public and private signals and the fast decline of private sentiments may have predicted or even contributed to the subsequent market crisis. This, to some extent, reflects the importance and power of narratives in a crisis.

### 3.3 Precisions

As Figure 3 corresponds to $y$ and $x$ in the model, we now turn to their precisions, $\tau$ and $\sigma$. We calculate the standard deviations of daily sentiments weighted by convictions to represent the noise levels of public and private signals, respectively. The standard deviations capture the extent of disagreement among different news sources and investors' sentiments. Figure 4 shows the results. As above, we only plot trading days to eliminate holiday effects. The blue lines, representing noise in comment sentiments, are higher than the red lines, which plot noise in news sentiments. This aligns with our findings above and again implies a high $\gamma$.

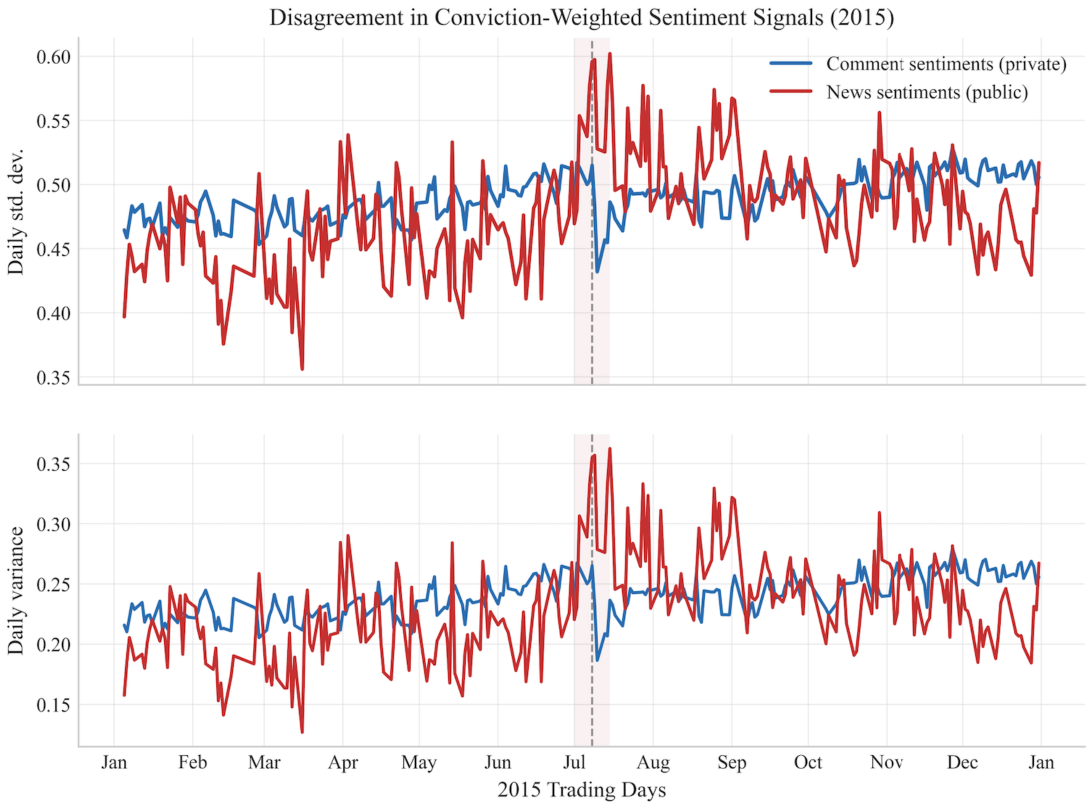


**Figure 4:** The standard deviation and variance of public and private signals

In addition, we observe another interesting phenomenon in Figure 4. At the beginning of July, right after the market crash, disagreement in news sentiments surged. This may reflect the divergence of authoritative media and other media. Just as the gap between comments and news sentiments aforementioned, the divergence between authoritative narratives and others may also be the consequence of manipulated counter-cyclical efforts by authoritative media to stabilize the market.

# 4 Analyzing the Model

We have already done comparative statics in section 2. In this section, we calibrate the model using sentiments extracted from textual data. We also perform counterfactual manipulation level analysis to show the existance of optimal manipulation level. Based on our results, we give our policy implications. It is useful to focus my analysis on the period from March to the end of Augest, when the market experienced a significant boom and bust cycle. The SSE Index rose from right below 3,500 points at the beginning of March to its peak of 5,166 on June 12, and then crashed to 3,507 on July 8, followed by a volidating rebound to around 4,000, but finally fell back to 2,927 on Augest 26 (See Figure 1). This period captures the most dramatic market movements in 2015, providing a rich context to study the effects of narratives and skepticism.

## 4.1 Model Parameters

We first calibrate parameters in the Tacitus Trap model under the *all-news* specification. There are three parameters to be calibrated in the endogenous skepticism model. One is the significance level $\alpha$, which governs investors' propensity to reject the public signal and therefore determines the portion of Believers and Skeptics in equilibrium. Another is the cost-of-holding parameter (intercept term in the empirical payoff specification), which shifts the marginal holding condition and helps the model fit the dynamics of holding behavior. The third is a rescale factor applied to the empirically measured noise terms in the behavioral / threshold calibration, which maps the sentiment-based noise measures into the model scale used to solve the threshold system. Notice that $\Delta$ is exogenous in the Tacitus Trap calibration and is therefore not calibrated in the threshold equations. In the baseline all-news calibration, the selected parameter values are $\alpha = 0.2$, cost-of-holding parameter $= 0.55$, and rescale factor $= 0.65$. This behavioral rescale factor is distinct from the separate technical rescaling used later for computing the baseline $\gamma$ and $\zeta$ series.

We set $\alpha = 0.2$, which means investors are willing to accept a 20% chance of making a wrong decision(Type I error) when testing the official narrative. As the normal distributions of noises are assumed in the model, that is around 0.84 standard deviations away from the mean. This value is consistent with the idea that during a crisis, investors may be more cautious and skeptical.

For the intercept term of the payoff function, recall that the benchmark payoff function is given by $\theta + l - 1$ in the model. In the calibration, however, we treat the cost-of-holding intercept as a parameter to be calibrated jointly with $\alpha$ and the behavioral noise rescale factor. Under the baseline all-news specification, the selected value is 0.55. This parameter shifts the holding threshold and helps the model fit the observed levels and dynamics of investor holding behavior. The rescale factor, selected as 0.65, is applied to the empirically measured public and private noise terms before solving the threshold equations.

## 4.2 Predicting the Crisis

Figure 5 shows the calibrated Tacitus Trap threshold $\kappa_T$ based on the aforementioned parameters and real world data. The light green and red points represent private signals $x_i$ each day (non-trading days are left blank), while the grass green line plots the daily Tacitus Trap threshold $\kappa_T$, and the dark green line is its five-day moving average. It can be seen that before the market crash, the threshold stayed at $\bar{x}(y)$ almost everyday, indicating that Believers chose to hold and Skeptics chose to run, roughly.

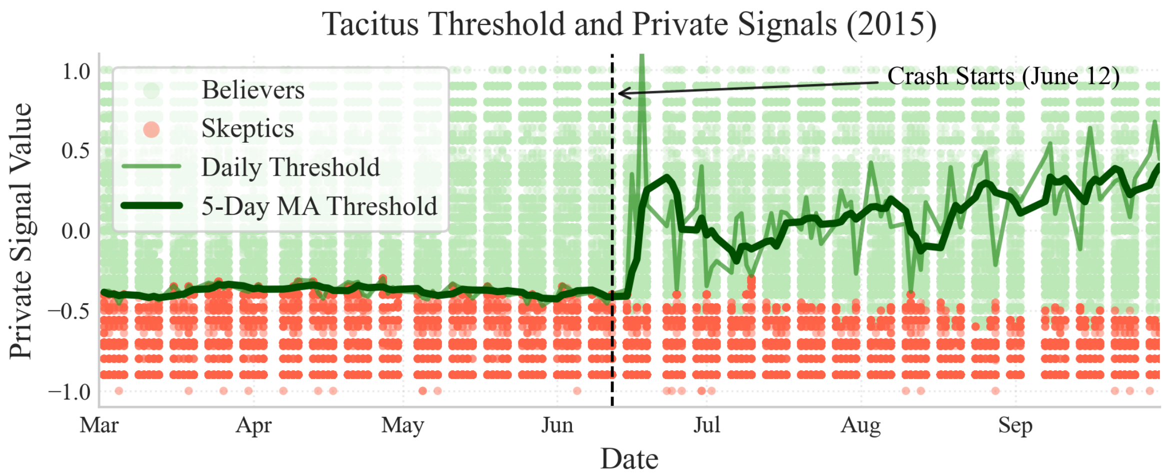


**Figure 5:** Calibrated equilibrium private signal threshold $\kappa^*$ with private signals and five-day moving average. Points above the threshold will hold, while points below the threshold will run. Green points are believers, red points are skeptics.

However, after the crash, the threshold jumped to $\kappa_S$ which is far higher than in normal times and is also higher than most private signals. This means that a large proportion of investors chose to run, regardless of the optimistic public narratives. This pattern aligns well with the Tacitus Trap mechanism, where public signals lose their coordinating power when skepticism rises sharply. In addition, as shown in Figure 1, the market crash can be separated into two phases: i) a sharp decline from June 12 to July 8, and ii) another sharp decline from August 17 (SSE Index: 3994) to August 26. Our calibrated Tacitus Trap threshold in Figure 5 also reflects this two-phase crash pattern, with the threshold dramatically high in the two phases.

To further illustrate the result, Figure 6 plots the calibrated proportions of investors holding $L$ (Panel A) and the proportion of Skeptics $\eta$ (Panel B) in the crisis. Again, the light lines are daily proportions and the dark lines are their five-day moving averages. One can see that before the crash, the proportion of holding was quite high, averagely around $80\%$, while the proportion of Skeptics was low, around $15\%$. After the crash, however, the proportion of holding dropped to below $40\%$ in June and below $20\%$ in August, while the proportion of Skeptics surged to around $25\%$. This shift is consistent with the previous observation of the Tacitus Trap threshold, and the model.

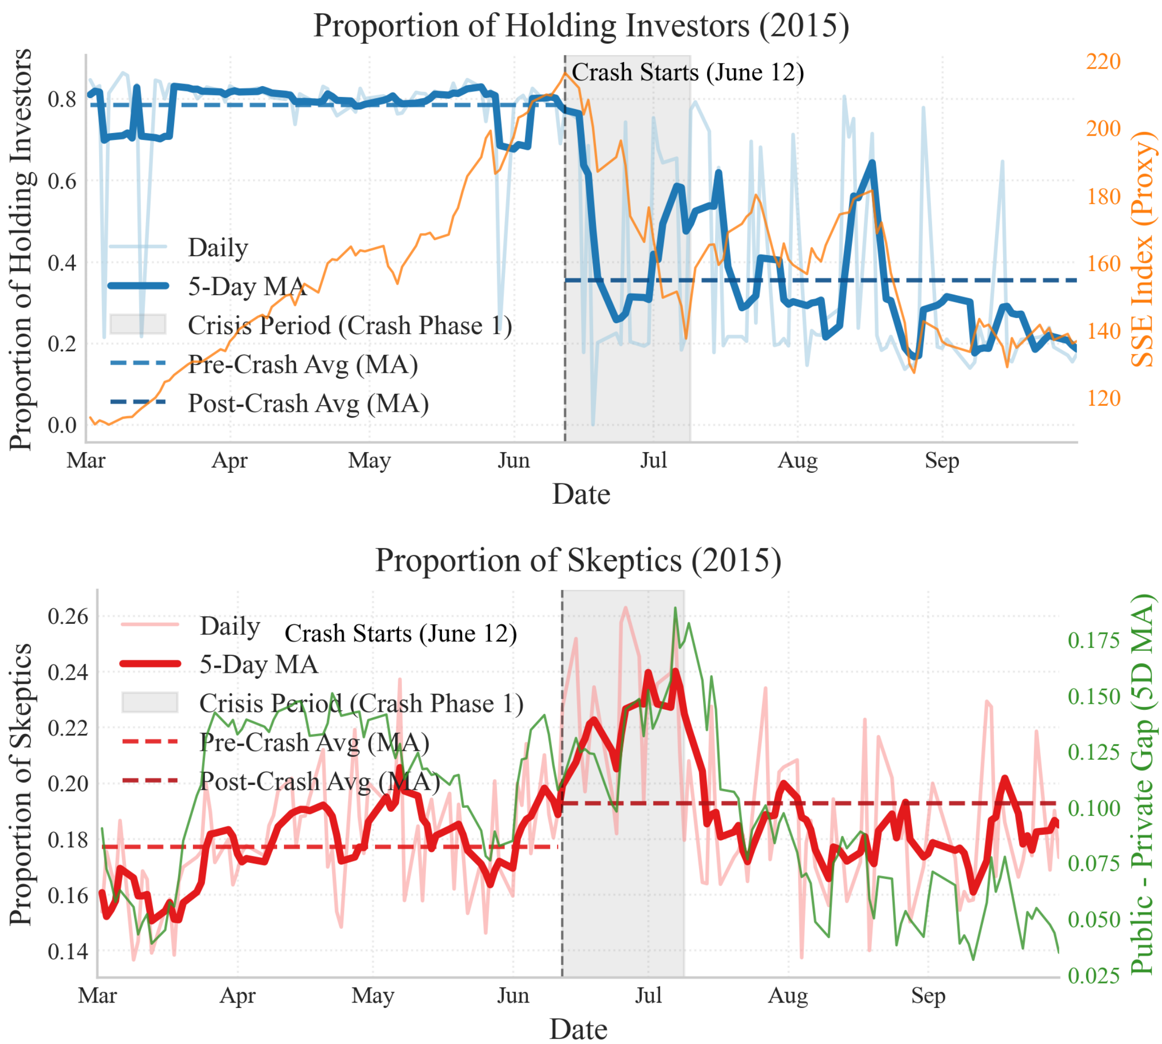


**Figure 6:** Calibrated proportions of investors holding $L$ (Panel A) and Skeptics $\eta$ (Panel B) during the crisis.

Diving into Panel A of Figure 6, where the SSE Index is also plotted, it is interesting to note that the sharp decline in the proportion of holding occurred at the same time when the SSE Index crashed, or even before that. For instance, right around June 12, there are two days when the calibrated proportion of holding dropped to zero, exactly before a decline of around 500 points in the Index. This implies that narratives and people's beliefs may contribute to the market crash, rather than simply reflecting the crash.

It is also interesting to observe in Panel B of Figure 6, where the gap between public and private signals is also plotted. One can easily find a large increase of the gap during the crash, implying strong evidence of public signal manipulation. Panel B tells the story of the Tacitus Trap. As the gap widened, the proportion of Skeptics surged, leading to a collapse of coordination effects (and fundamental effects as well) of public narratives, ultimately contributing to the market crash. It is the diminished "coordination effect" and the surging "Tacitus Trap effect" that disrupt the balanced trade-off, leading to a self-fulfilling crash

In conclusion, the calibrated Tacitus Trap model successfully captures the dynamics of investor behavior and macro market results during the 2015 Chinese stock market crisis. The model addresses the important role of narratives in a crisis. Moreover, it shows how endogenous skepticism can lead to a breakdown of fundamental and coordination effects of public signals.

## 4.3 Analyzing Coordination Effects

We calculate the empirical relative precision of public over private signals $\gamma$ and the empirical publicity multiplier $\zeta$ under the assumptions and formulas in the baseline model, i.e., without considering endogenous skepticism.

Notice that we have to technically rescale the standard deviations of both signals by 1.358 times in order to match the model, due to the restriction that $\gamma < 2\pi \approx 6.283$. In the empirical data, there are some days when $\gamma$ exceeds this upper bound, leading to negative values of $\zeta$. A very high $\gamma$ indicates that public signals are perceived as extremely concentrated compared to private signals, which in some cases (e.g., the launch of particular policies) is realistic. However, the rescaling will not affect the order of magnitudes of $\gamma$ across different days, thus preserving the relative comparison of publicity multipliers $\zeta$.

Figures 7 and 8 plot the daily values of $\gamma$ and $\zeta$ in 2015, respectively, along with their five-day moving averages to smooth trading-day effects. It can be seen that $\gamma$ is higher on the first half of the year, before the crash, and then decreases afterwards. This coincides with the observation that disagreements in news sentiments increased after the crash. For $\zeta$, on the second half of the year, the coordination effects of public signals weakened significantly, indicating that public signals may have lost their power to coordinate investors during the crisis period.

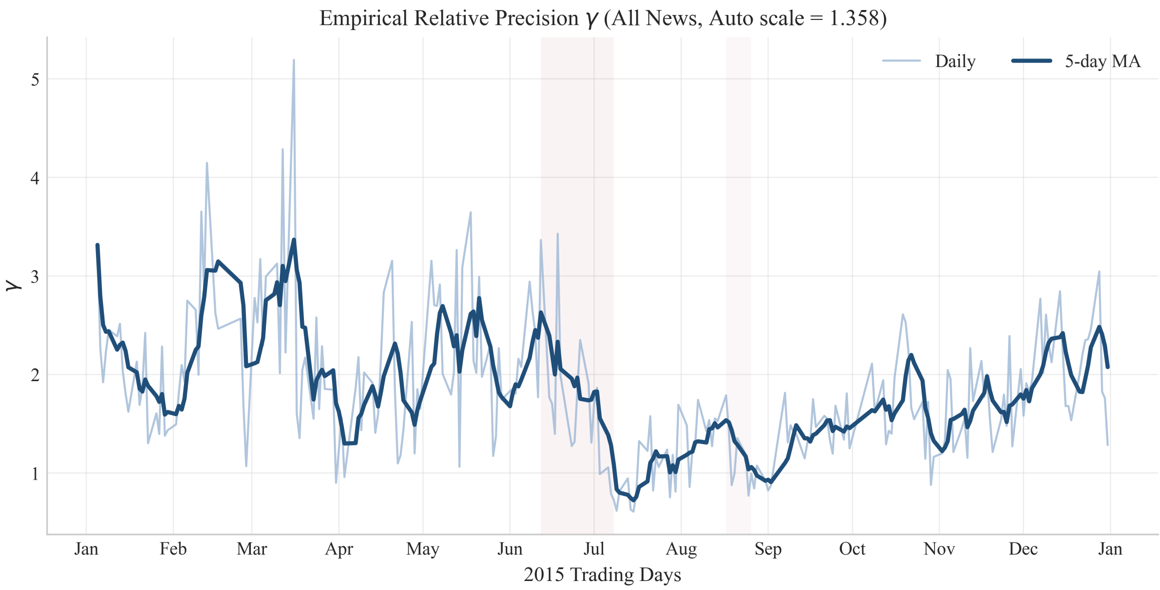


**Figure 7:** Empirical relative precision of public over private signals($\gamma$)

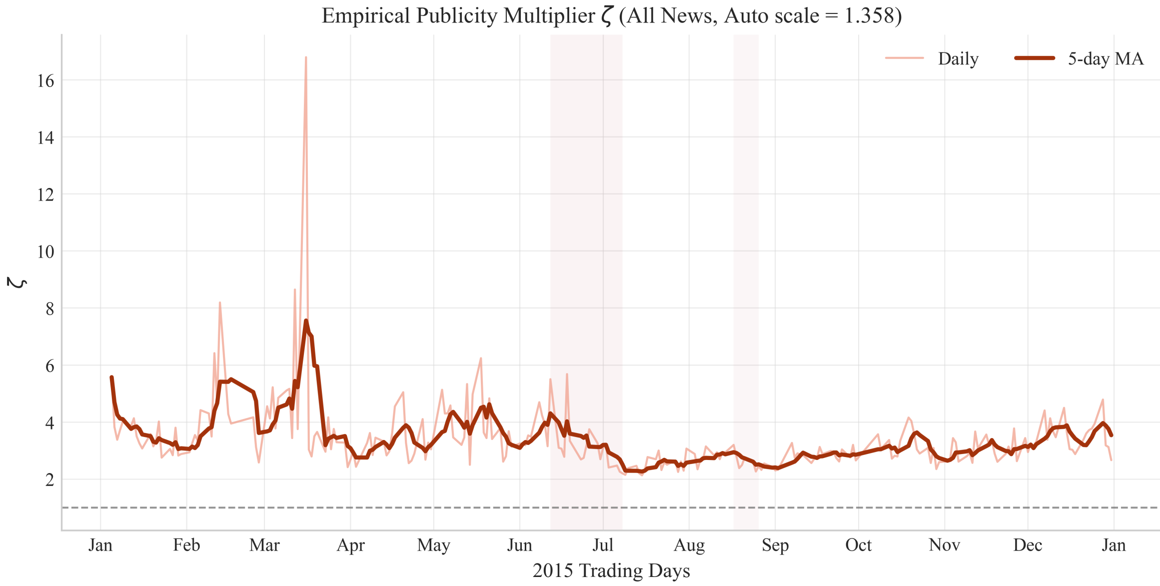


**Figure 8:** Empirical publicity multiplier($\zeta$), capturing the coordination effect of public signals.

#### 4.3.1 Outliers in Publicity Multiplier

In Figure 8, there are several spikes in the publicity multiplier $\zeta$. These outliers provide an excellent opportunity to validate the model's prediction that public signals can serve as powerful coordination devices.

Table 4 lists the top five dates with the highest $\zeta$ values, along with key public events around those dates. These dates broadly coincide with major policy-signaling windows and salient public events, including the mid-March policy communication window (surrounding the CSRC Chairman's "Reform Bull" statement on March 10) and the release window of the "Made in China 2025" plan on May 18. Such events likely served as powerful coordination devices for investors, contributing to elevated publicity multipliers on those days. This pattern increases confidence in our text-based measurement and provides qualitative support for the model's mechanism.

**Table 4:** Top 5 Dates with Highest Multiplier ($\zeta$) and Corresponding Events. The table lists the five trading days with the highest realized publicity multiplier effect ($\zeta$) in 2015. High $\zeta$ values coincide with major authoritative policy announcements, validating the model's prediction that public signals serve as powerful coordination devices during grand narrative events.

| **Date** | $\zeta$ | **Key Public Signals** |
|---|---|---|
| Mar. 16 | 16.80 | Post–Two Sessions policy-signaling window |
| Mar. 12 | 8.65 | NPC press conference on financial reform and development |
| Feb. 13 | 8.19 | Pre-Lunar-New-Year liquidity and policy-expectation window |
| Feb. 11 | 6.42 | First week of SSE 50ETF options trading |
| May. 18 | 6.24 | "Made in China 2025"policy release window |

## 4.4 The Optimal Manipulation

Using the calibrated all-news parameters ($\alpha = 0.2$, cost-of-holding $= 0.55$, rescale factor $= 0.65$), we conduct a counterfactual exercise that varies only the public signal $y$. The purpose is not to recover a realistic day-by-day policy rule, but to test the model's mechanism and quantify the trade-off between coodination effect and extreme tacirus trap.

We first compute, for each trading day $t$, the counterfactual shift $\Delta_y^*(t)$ that maximizes the same-day holding ratio. Equivalently, we define the gap between the realized public signal and the daily holding-maximizing threshold, $g_t \equiv y_t - y^*_{t,\text{hold}}$.

A negative $g_t$ means that the realized public signal is below the threshold that would maximize the holding ratio on that day. In this region, a moderate upward shift in public narratives would increase the aggregate holding ratio because the coordination effect of public signals dominates the credibility loss. However, this does not imply that "more positive is always better": once the manipulation becomes too large, the gap between public and private signals widens, skepticism rises, and the Tacitus Trap effect starts to dominate, so the holding ratio declines. The counterfactual exercise therefore identifies a trade-off between coordination gains (at low to moderate manipulation levels) and credibility losses (at high manipulation levels).

Figure 9 plots this gap during the crash episodes. The gap is close to zero or slightly positive around the crash onset, but becomes predominantly negative afterwards, especially during the second crash phase (August 17–26), indicating that the required signal support rises sharply as skepticism intensifies.

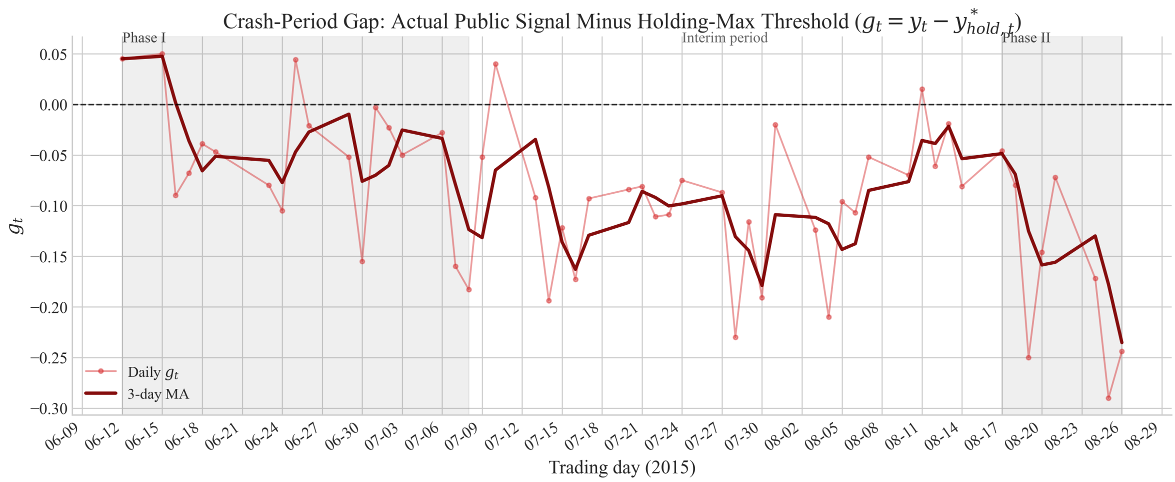


**Figure 9:** $g_t \equiv y_t - y^*_{t,\text{hold}}$ during the crash episodes. This shows the gap between realized public signal and the public signal under optimal manipulation.

However, as an ex-ante public media agent, we have no idea on what will happen precisely every day. So, in order to obtain a policy-relevant benchmark, we then incorporate a uniform counterfactual manipulation $\Delta$. throghout the whole crisis window and then compute the average holding ratio.

Figure 10 shows a clear inverted-U relationship: the average holding ratio increases for small values of $\Delta$, reaches a maximum at an intermediate level, and then declines as larger manipulation reduces credibility. Which dynamically shows the trade-off of coodination effect and Tacitus Trap effect.

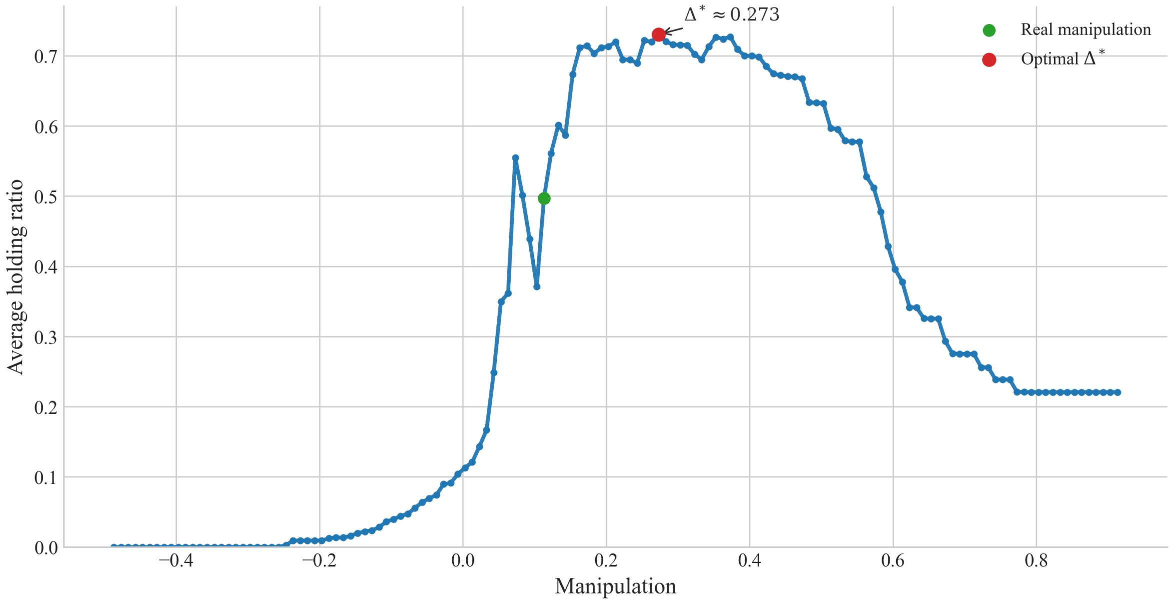


**Figure 10:** Manipulation level v.s. average holding ratio. As manipulation level increases, average holding ratio first increases and then decreases. $\Delta^*$ is the optimal manipulation level and the green point is the real manipulation level.

For the crisis-combined window, the simulated optimum is around $\Delta^* = 0.273$, whereas the real manipulation is lower (around $\Delta_{\text{real}} \approx 0.097$). These results provide quantitative support for the model's central prediction of a non-monotonic effect of public-signal manipulation.

### 4.4.1 Policy Implication

Following the analysis in the above section, we suggest that public media should release more optimistic narratives. Besides, the result suggests a more subtle principle: in crisis periods, public media should be willing to provide narrative support, but such support must be cautious.

Our findings show that when the realized public signal is below the holding-maximizing threshold, a moderate increase in public narrative support can improve market stability, because the coordination effect of public signals still dominates the Tacitus Trap effect. In this sense, public media should not be completely passive during a crisis. Timely and supportive narratives may help stabilize expectations and prevent panic from becoming self-fulfilling.

However, from an ex ante perspective, policymakers and public media cannot know the exact optimal manipulation level in real time. They must act under uncertainty, and they must also take into account the potential damage from the Tacitus Trap once credibility is weakened. This concern is even more important in the presence of explicit market frictions, such as price-limit constraints (e.g., limit-down episodes), which can amplify panic selling and make credibility breakdown more costly.

Therefore, the practical implication is cautious encouragement rather than aggressive persuasion. Public media should provide support that is moderate, consistent with observable market conditions, and backed by concrete policy actions and verifiable information. Once the gap between public and private signals becomes too large and skepticism rises, communication should shift from stronger encouragement to more transparent risk disclosure and policy clarification, so as to avoid pushing the market into a Tacitus Trap equilibrium.

# 5 Conclusion

This paper uses a global games model with endogenous skepticism to analyze the role of narratives in financial market crises. We find there exists an optimal manipulation level of public news during a crisis. The model also highlights the roles of private and public signals in driving investor behavior and market returns, as well as showing that public signals have a strong coordination effect in normal times. The model is calibrated using textual data from news articles and social media comments during the 2015 Chinese stock market crash. We use the novel Qwen2.5-14B-Instruct-AWQ model in local environment to extract sentiments from news and comments, which can capture not only sentiments in individual words but also the overall narratives conveyed by the texts. Our model and data fit well with reality, proving the reliability of the LLM tool and the model.

However, there are also several limitations in this paper. First, LLM-based scoring may contain model-specific biases; future work can leverage multiple models, calibration sets, or human audits to further validate measurement invariance, which will definitely solve the problem of alignment. Second, our model simplifies the mechanism of crisis to better capture the influence of narrative manipulation, which may be weak while facing casual empirical tests. Finally, our current setting is static, which means investors cannot learn from past manipulations. In the future, we may shift to a dynamic setting to be more realistic and gain more insights.

# References


**An, Li, Dong Lou, and Donghui Shi**, "Wealth redistribution in bubbles and crashes," *Journal of Monetary Economics*, 2022, *126*, 134–153.

**Angeletos, George-Marios and Alessandro Pavan**, "Efficient use of information and social value of information," *Econometrica*, 2007, *75* (4), 1103–1142.

**__ , Christian Hellwig, and Alessandro Pavan**, "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks," *Econometrica*, 2007, *75* (3), 711–756.

**Antweiler, Werner and Murray Z Frank**, "Is all that talk just noise? The information content of internet stock message boards," *The Journal of Finance*, 2004, *59* (3), 1259–1294.

**Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer**, "Extrapolation and bubbles," *Journal of Financial Economics*, 2018, *129* (2), 203–227.

**Bian, Jiangze, Zhi Da, Zhiguo He, Dong Lou, Kelly Shue, and Hao Zhou**, "The Drivers and Implications of Retail Margin Trading," Working Paper August 2025. Draft version. This paper subsumes previous working papers: "Leverage-Induced Fire Sales and Stock Market Crashes" and "Leverage Networks and Market Contagion".

**Bond, Philip and Itay Goldstein**, "Government Intervention and Information Aggregation by Prices," *The Journal of Finance*, 2015, *70* (6), 2777–2812.

**Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, "Diagnostic expectations and credit cycles," *The Journal of Finance*, 2018, *73* (1), 199–227.

**Brunnermeier, Markus K, Michael Sockin, and Wei Xiong**, "China's Model of Managing the Financial System," *The Review of Economic Studies*, 2022, *89* (6), 3115–3153.

**Bybee, Leland, Bryan Kelly, and Yinan Su**, "Narrative asset pricing: Interpretable systematic risk factors from news text," *The Review of Financial Studies*, 2023, *36* (12), 4759–4787.

**Carlsson, Hans and Eric van Damme**, "Global Games and Equilibrium Selection," *Econometrica*, 1993, *61* (5), 989–1018.

**Diamond, Douglas W and Philip H Dybvig**, "Bank runs, deposit insurance, and liquidity," *Journal of Political Economy*, 1983, *91* (3), 401–419.

**Edmond, Chris**, "Information manipulation, coordination, and regime change," *The Review of Economic Studies*, 2013, *80* (4), 1422–1458.

**Eliaz, Kfir and Ran Spiegler**, "A model of competing narratives," *American Economic Review*, 2020, *110* (12), 3786–3816.

**Fang, Hanming, Ming Li, and Guangli Lu**, "Decoding China's Industrial Policies," NBER Working Paper 33814, National Bureau of Economic Research 2025.

**Frankel, David M., Stephen Morris, and Ady Pauzner**, "Equilibrium selection in global games with strategic complementarities," *Journal of Economic Theory*, 2003, *108*, 1–44.

**García, Diego**, "Sentiment during recessions," *The Journal of Finance*, 2013, *68* (3), 1267–1300.

**Gennotte, Gerard and Hayne Leland**, "Market liquidity, hedging, and crashes," *The American Economic Review*, 1990, *80* (5), 999–1021.

**Gentzkow, Matthew and Jesse M Shapiro**, "Media bias and reputation," *Journal of Political Economy*, 2006, *114* (2), 280–316.

**Goetzmann, William N, Dasol Kim, and Robert J Shiller**, "Crash Narratives," Working Paper 30195, National Bureau of Economic Research 2024. Revised November 2024.

**Goldstein, Itay and Ady Pauzner**, "Demand-deposit contracts and the probability of bank runs," *The Journal of Finance*, 2005, *60* (3), 1293–1327.

**Greenwood, Robin and Andrei Shleifer**, "Expectations of returns and expected returns," *The Review of Financial Studies*, 2014, *27* (3), 714–746.

**He, Zhiguo and Wei Xiong**, "Dynamic Debt Runs," *The Review of Financial Studies*, 2012, *25* (6).

**Lopez-Lira, Alejandro and Yuehua Tang**, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," *arXiv preprint arXiv:2304.07619*, 2025. Revised October 2025.

**Loughran, Tim and Bill McDonald**, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

**Manela, Asaf and Alan Moreira**, "News implied volatility and disaster concerns," *Journal of Financial Economics*, 2017, *123* (2), 137–162.

**Morris, Stephen and Hyun Song Shin**, "Global games: Theory and applications," in "Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress," Vol. 1, Cambridge University Press, 2003, pp. 56–114.

**Obstfeld, Maurice**, "Models of currency crises with self-fulfilling features," *European Economic Review*, 1996, *40* (3-5), 1037–1047.

**Ortoleva, Pietro**, "Modeling the change of paradigm: Non-Bayesian reactions to unexpected news," *American Economic Review*, 2012, *102* (6), 2410–2436.

**Rabin, Matthew and Joel L Schrag**, "First impressions matter: A model of confirmatory bias," *The Quarterly Journal of Economics*, 1999, *114* (1), 37–82.

**Shiller, Robert J**, "Narrative economics," *American Economic Review*, 2017, *107* (4), 967–1004.

**Tetlock, Paul C**, "Giving content to investor sentiment: The role of media in the stock market," *The Journal of Finance*, 2007, *62* (3), 1139–1168.

# Appendix

# A Proof of Proposition 1

This appendix provides the formal proof of Proposition 1, establishing existence and uniqueness of the signal threshold equilibrium.

## A.1 Setup and Notation

Recall the value function for an investor with private signal $x$:

$$V(x) = \begin{cases} x + \mathbb{E}[L(\theta, y)|x] - 1, & \text{if } x < \bar{x}(y) \quad \text{(Skeptic Region)} \\ [\rho x + (1-\rho)y] + \mathbb{E}[L(\theta, y)|x, y] - 1, & \text{if } x \geq \bar{x}(y) \quad \text{(Believer Region)}, \end{cases} \tag{24}$$

where $\rho = \frac{\tau^2}{\sigma^2+\tau^2}$, $\bar{x}(y) = y - z_\alpha \sqrt{\sigma^2 + \tau^2}$, and $L(\theta, y)$ is the aggregate holding mass. In a threshold equilibrium where all investors with $x_i > \kappa$ hold, we have

$$L(\theta, y) = 1 - \Phi\left(\frac{\kappa - \theta}{\sigma}\right), \tag{25}$$

which is strictly increasing in $\theta$ for any finite $\kappa$.

We need to show that there exists a unique $\kappa^*$ such that $V(\kappa^*) = 0$.

## A.2 Monotonicity in the Skeptic Region

For $x < \bar{x}(y)$, a Skeptic's posterior belief is $\theta|x \sim N(x, \sigma^2)$. The expected holding mass is

$$\mathbb{E}[L(\theta, y)|x] = \int_{-\infty}^{\infty} L(\theta, y) \frac{1}{\sigma} \phi\left(\frac{\theta - x}{\sigma}\right) d\theta. \tag{26}$$

Using the threshold form (25) and the normal convolution identity,[1] this evaluates to

$$\mathbb{E}[L|x] = 1 - \Phi\left(\frac{\kappa - x}{\sqrt{\sigma^2 + \sigma^2}}\right) = \Phi\left(\frac{x - \kappa}{\sigma\sqrt{2}}\right). \tag{27}$$

Differentiating with respect to $x$:

$$\frac{dV^S}{dx} = 1 + \frac{1}{\sigma\sqrt{2}} \phi\left(\frac{x - \kappa}{\sigma\sqrt{2}}\right) > 1 > 0. \tag{28}$$

Hence $V^S$ is strictly increasing in the Skeptic region.

[1] For independent $Z_1 \sim N(\mu_1, \sigma_1^2)$ and $Z_2 \sim N(0, \sigma_2^2)$, $\int \Phi\left(\frac{a-z}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{z-\mu_1}{\sigma_1}\right) dz = \Phi\left(\frac{a-\mu_1}{\sqrt{\sigma_1^2+\sigma_2^2}}\right)$.

## A.3 Monotonicity in the Believer Region

For $x \geq \bar{x}(y)$, a Believer's posterior belief is $\theta|x, y \sim N(\overline{\theta}, \tilde{\sigma}^2)$, where $\overline{\theta} = \rho x + (1-\rho)y$ and $\tilde{\sigma}^2 = \frac{\sigma^2\tau^2}{\sigma^2+\tau^2}$. The value function is

$$V^B(x) = \overline{\theta} + \mathbb{E}[L(\theta, y)|\overline{\theta}] - 1. \tag{29}$$

Using (25) and the same normal convolution identity, the expected holding mass under the Believer's posterior is

$$\mathbb{E}[L|\overline{\theta}] = 1 - \Phi\left(\frac{\kappa - \overline{\theta}}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) = \Phi\left(\frac{\overline{\theta} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right). \tag{30}$$

This is the convolution of two normal CDFs: the holding mass $L(\theta, y) = \Phi\left(\frac{\theta-\kappa}{\sigma}\right)$ integrated against the Believer's posterior density $\frac{1}{\tilde{\sigma}}\phi\left(\frac{\theta-\overline{\theta}}{\tilde{\sigma}}\right)$. The resulting composite variance is $\sigma^2 + \tilde{\sigma}^2$.

Differentiating (30) with respect to $\overline{\theta}$:

$$\frac{\partial \mathbb{E}[L]}{\partial \overline{\theta}} = \frac{1}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\,\phi\left(\frac{\overline{\theta} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) \geq 0. \tag{31}$$

This derivative is *non-negative* for any candidate threshold $\kappa$: a higher posterior mean shifts the Believer's posterior to the right, increasing the expected proportion of investors who hold. Using $\frac{d\overline{\theta}}{dx} = \rho > 0$:

$$\frac{dV^B}{dx} = \rho\left(1 + \frac{\partial \mathbb{E}[L]}{\partial \overline{\theta}}\right) \geq \rho > 0. \tag{32}$$

Hence $V^B$ is strictly increasing in the Believer region.

## A.4 Boundary Behavior

**As $x \to -\infty$:** In the Skeptic region, $V^S(x) = x + \mathbb{E}[L|x] - 1$. Since $\mathbb{E}[L|x] \in [0, 1]$ is bounded, $V^S(x) \to -\infty$ as $x \to -\infty$.

**As $x \to +\infty$:** In the Believer region, $V^B(x) = \rho x + (1-\rho)y + \mathbb{E}[L|x, y] - 1$. Since $\rho > 0$ and $\mathbb{E}[L]$ is bounded, $V^B(x) \to +\infty$ as $x \to +\infty$.

## A.5 Continuity and the Jump at the Cutoff

At $x = \bar{x}(y)$, the value function may be discontinuous because an investor who crosses the skepticism cutoff switches from discarding to incorporating the public signal. The left-hand limit (Skeptic side) is

$$V^S(\bar{x}(y)^-) = \bar{x}(y) + \Phi\left(\frac{\bar{x}(y) - \kappa}{\sigma\sqrt{2}}\right) - 1, \tag{33}$$

and the right-hand limit (Believer side) is

$$V^B(\bar{x}(y)^+) = \underbrace{\rho\bar{x}(y) + (1-\rho)y}_{\overline{\theta}_{\bar{x}}} + \Phi\left(\frac{\overline{\theta}_{\bar{x}} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) - 1, \tag{34}$$

where $\overline{\theta}_{\bar{x}} = \rho\bar{x}(y) + (1-\rho)y$.

The jump in the value function is

$$V^B(\bar{x}(y)^+) - V^S(\bar{x}(y)^-) = \underbrace{(1-\rho)(y - \bar{x}(y))}_{\text{posterior mean increase}} + \underbrace{\Phi\left(\frac{\overline{\theta}_{\bar{x}} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) - \Phi\left(\frac{\bar{x}(y) - \kappa}{\sigma\sqrt{2}}\right)}_{\text{change in expected coordination}}. \tag{35}$$

The first term is strictly positive since $y - \bar{x}(y) = z_\alpha\sqrt{\sigma^2 + \tau^2} > 0$ and $(1-\rho) > 0$. For the second term, note that $\overline{\theta}_{\bar{x}} = \bar{x}(y) + (1-\rho)(y - \bar{x}(y)) > \bar{x}(y)$. So

$$\frac{\overline{\theta}_{\bar{x}} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}} > \frac{\overline{\theta}_{\bar{x}} - \kappa}{\sigma\sqrt{2}} > \frac{\bar{x}(y) - \kappa}{\sigma\sqrt{2}}.$$

The first inequality is given by $\sqrt{\sigma^2 + \tilde{\sigma}^2} = \sqrt{(1 + \frac{\tau^2}{\sigma^2+\tau^2})\sigma^2} < \sigma\sqrt{2}$. Both effects increase the CDF argument, making the second term non-negative. Therefore, $V^B(\bar{x}^+) > V^S(\bar{x}^-)$: the jump is strictly upward.

## A.6 Existence and Uniqueness

Combining the results from the previous subsections:

1. $V^S(x)$ is strictly increasing for $x < \bar{x}(y)$, and $V^B(x)$ is strictly increasing for $x \geq \bar{x}(y)$.
2. $V(x) \to -\infty$ as $x \to -\infty$ and $V(x) \to +\infty$ as $x \to +\infty$.
3. There is a strictly upward jump at $x = \bar{x}(y)$: $V^B(\bar{x}(y)^+) > V^S(\bar{x}(y)^-)$.

Properties (1) and (2) together with the intermediate value theorem guarantee the existence of at least one $\kappa^*$ with $V(\kappa^*) = 0$. Uniqueness follows from strict monotonicity within each region combined with the upward jump: the value function cannot cross zero more than once. Formally:

**Case Analysis:**

- *Case 1:* $V^S(\bar{x}(y)^-) < 0$ and $V^B(\bar{x}(y)^+) > 0$. Then $V$ crosses zero exactly at the boundary. The equilibrium threshold is $\kappa^* = \bar{x}(y)$.
- *Case 2:* $V^S(\bar{x}(y)^-) \geq 0$. Since $V^S$ is strictly increasing with $V^S(-\infty) = -\infty$, there exists a unique $\kappa^* < \bar{x}(y)$ in the Skeptic region satisfying $V^S(\kappa^*) = 0$. By the upward jump, $V^B(x) > 0$ for all $x \geq \bar{x}(y)$, so there is no additional zero in the Believer region. We denote $\kappa^* = \kappa^*_S$.
- *Case 3:* $V^B(\bar{x}(y)^+) \leq 0$. Since $V^S(\bar{x}(y)^-) < V^B(\bar{x}(y)^+) \leq 0$, and $V^S$ is strictly increasing from $-\infty$, $V^S(x) < 0$ for all $x \leq \bar{x}(y)$. Then $V^B$ is strictly increasing with $V^B(+\infty) = +\infty$, so there exists a unique $\kappa^* > \bar{x}(y)$ in the Believer region satisfying $V^B(\kappa^*) = 0$. We denote $\kappa^* = \kappa^*_B$.

## A.7 Relationship Between $\kappa_S$ and $\kappa_B$

We establish that $\kappa_S^* > \kappa_B^*$. Consider the two indifference conditions evaluated at a common signal $x$:

$$V^S(x) = x + \Phi\left(\frac{x - \kappa_S^*}{\sigma\sqrt{2}}\right) - 1, \tag{36}$$

$$V^B(x) = [\rho x + (1-\rho)y] + \Phi\left(\frac{\overline{\theta}(x) - \kappa_B^*}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) - 1. \tag{37}$$

The Believer's value exceeds the Skeptic's value at any common $x$ in the region where $y > x$ (which includes the neighborhood of the skepticism cutoff $\bar{x}(y)$, since $y > \bar{x}(y)$) for two reasons: (i) the Believer's posterior mean $\rho x + (1-\rho)y > x$ is higher, and (ii) the Believer's more precise posterior (smaller variance $\tilde{\sigma}^2 < \sigma^2$) leads to tighter concentration around the posterior mean, yielding (weakly) higher expected coordination.

Formally, suppose the marginal Skeptic is indifferent: $V^S(\kappa_S^*) = 0$. Then for the Believer evaluated at the same signal $x = \kappa_S^*$:

$$V^B(\kappa_S^*) = [\rho\kappa_S^* + (1-\rho)y] + \mathbb{E}[L|\kappa_S^*, y] - 1 > \kappa_S^* + \mathbb{E}[L|\kappa_S^*] - 1 = V^S(\kappa_S^*) = 0, \tag{38}$$

where the inequality uses $\rho\kappa_S^* + (1-\rho)y > \kappa_S^*$ (since $y > \kappa_S^*$ when $\kappa_S^* < \bar{x}(y) < y$) and $\mathbb{E}[L|\kappa_S^*, y] \geq \mathbb{E}[L|\kappa_S^*]$ (Believers have a higher posterior mean and more precise information). Since $V^B(\kappa_S^*) > 0$ and $V^B$ is strictly increasing, the Believer's threshold must satisfy $\kappa_B^* < \kappa_S^*$.

This completes the proof of Proposition 1. □

# B Proof of Proposition 3

Fix $\theta \in (0,1)$ and condition on $\epsilon^y = 0$, so that $y = \theta + \Delta$. The skepticism cutoff is $\bar{x}(\Delta) \equiv y - c = \theta + \Delta - c$, where $c = z_\alpha\sqrt{\sigma^2 + \tau^2} > 0$. The equilibrium signal threshold is $\kappa^*(\Delta)$, which is the unique solution to $V(\kappa^*; \Delta) = 0$ as established in Proposition 1. The aggregate holding mass is

$$\bar{L}(\theta; \Delta) = \left[\Phi\left(\frac{\bar{x}(\Delta) - \theta}{\sigma}\right) - \Phi\left(\frac{\kappa_S^*(\Delta) - \theta}{\sigma}\right)\right]^+ + 1 - \Phi\left(\frac{\max(\kappa_B^*(\Delta),\, \bar{x}(\Delta)) - \theta}{\sigma}\right). \tag{39}$$

The proof proceeds in four steps.

## B.1 Step 1: Regime Classification

As $\Delta$ varies, the relationship between $\bar{x}(\Delta)$ and the equilibrium thresholds $\kappa_B^*$, $\kappa_S^*$ determines three regimes.

**Regime A (No Skeptic Holds): $\bar{x}(\Delta) \leq \kappa_B^*(\Delta)$.** Since $\bar{x} \leq \kappa_B^* < \kappa_S^*$, any Skeptic (with $x_i < \bar{x}$) satisfies $x_i < \kappa_S^*$ and therefore does not hold. The only investors who hold are Believers with $x_i > \kappa_B^*$. Note that every Believer (with $x_i \geq \bar{x}$) faces the threshold $\kappa_B^*$, and since $\bar{x} \leq \kappa_B^*$, some Believers hold and some do not. The aggregate holding mass is determined solely by the Believer threshold:

$$\bar{L}(\theta; \Delta) = 1 - \Phi\left(\frac{\kappa_B^*(\Delta) - \theta}{\sigma}\right). \tag{40}$$

In this regime, Skeptics are present but contribute zero holding mass, so $\bar{L}$ takes the same form as in a pure Believer economy.

**Regime B (Mixed, All Believers Hold): $\kappa_B^*(\Delta) < \bar{x}(\Delta) \leq \kappa_S^*(\Delta)$.** Some investors are Skeptics, but no Skeptic holds: since $\bar{x} \leq \kappa_S^*$, every Skeptic (with $x_i < \bar{x} \leq \kappa_S^*$) falls below the Skeptic threshold and does not hold. Meanwhile, every Believer (with $x_i \geq \bar{x} > \kappa_B^*$) exceeds the Believer threshold and therefore holds. The aggregate holding mass equals the fraction of Believers:

$$\bar{L}(\theta; \Delta) = 1 - \Phi\left(\frac{\bar{x}(\Delta) - \theta}{\sigma}\right). \tag{41}$$

**Regime C (Both Types Hold): $\bar{x}(\Delta) > \kappa_S^*(\Delta)$.** Both types contribute to holding. Skeptics (with $x_i < \bar{x}$) hold if $x_i > \kappa_S^*$, contributing mass from $x_i \in (\kappa_S^*, \bar{x})$. Since $\bar{x} > \kappa_S^* > \kappa_B^*$, every Believer (with $x_i \geq \bar{x}$) exceeds the Believer threshold and holds. Summing:

$$\bar{L}(\theta; \Delta) = \Phi\left(\frac{\bar{x} - \theta}{\sigma}\right) - \Phi\left(\frac{\kappa_S^* - \theta}{\sigma}\right) + 1 - \Phi\left(\frac{\bar{x} - \theta}{\sigma}\right) = 1 - \Phi\left(\frac{\kappa_S^*(\Delta) - \theta}{\sigma}\right). \tag{42}$$

## B.2 Step 2: Part (i) —Holding Ratio is Increasing for Small Manipulation

Since $\phi(\cdot) \leq \frac{1}{\sqrt{2\pi}}$:

$$\frac{\partial \mathbb{E}[L]}{\partial \bar{\theta}} \leq \frac{1}{\sqrt{2\pi(\sigma^2 + \tilde{\sigma}^2)}} \leq \frac{1}{\tilde{\sigma}\sqrt{2\pi}} = \sqrt{\xi} < 1 \quad \text{under } \xi < 1. \tag{43}$$

Similarly, the derivative of $\mathbb{E}[L]$ with respect to the threshold $\kappa$ (holding $\bar{\theta}$ fixed) is

$$\frac{\partial \mathbb{E}[L]}{\partial \kappa} = -\frac{1}{\sqrt{\sigma^2 + \tilde{\sigma}^2}} \phi\left(\frac{\bar{\theta} - \kappa}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) \leq 0, \tag{44}$$

with $\left|\frac{\partial \mathbb{E}[L]}{\partial \kappa}\right| = \frac{\partial \mathbb{E}[L]}{\partial \bar{\theta}} \leq \frac{1}{\tilde{\sigma}\sqrt{2\pi}} < 1$ under $\xi < 1$.

By Assumption 1, $\bar{x}(0) = \theta - c < \kappa_B^*(0)$. Since $\bar{x}(\Delta) = \theta + \Delta - c$ is continuous and increasing in $\Delta$, and $\kappa_B^*(\Delta)$ is continuous in $\Delta$ (by the implicit function theorem applied to the indifference condition (10) under the regularity condition $\xi < 1$), there exists $\Delta_A > 0$ such that the economy is in Regime A for all $\Delta \in [0, \Delta_A)$.

In Regime A, $\bar{L}(\theta; \Delta) = 1 - \Phi\left(\frac{\kappa_B^*(\Delta) - \theta}{\sigma}\right)$, so $\bar{L}$ is increasing in $\Delta$ if and only if $\kappa_B^*(\Delta)$ is decreasing in $\Delta$.

**Claim:** In Regime A, $\frac{d\kappa_B^*}{d\Delta} < 0$.

*Proof.* In Regime A, no Skeptic holds, so the aggregate holding mass takes the same form as in a standard Morris and Shin (2003) global game with public signal $y = \theta + \Delta$: $L(\theta, y) = 1 - \Phi\left(\frac{\kappa_B^* - \theta}{\sigma}\right)$. We show $\frac{d\kappa_B^*}{d\Delta} < 0$ using monotone comparative statics, which does not require a contraction condition.

*Step (a): Best-response threshold decreases in $y$.* Fix any candidate threshold $k$ used by all other investors. The marginal Believer with signal $x$ has value function

$$V^B(x; k, y) = [\rho x + (1-\rho)y] + \Phi\left(\frac{\rho x + (1-\rho)y - k}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right) - 1. \tag{45}$$

Both the first term and the argument of $\Phi$ are strictly increasing in $y$ (for fixed $x$ and $k$). Therefore $V^B(x; k, y)$ is strictly increasing in $y$. The best-response threshold $BR(k; y)$, defined by $V^B(BR(k; y); k, y) = 0$, is strictly decreasing in $y$: when $y$ increases, the zero crossing of $V^B(\cdot; k, y)$ shifts to the left since $V^B$ has increased everywhere.

*Step (b): Best-response threshold increases in k (strategic complements).* For fixed $y$, increasing $k$ reduces $L(\theta, y)$ at every $\theta$, which decreases $\mathbb{E}[L|\bar{\theta}]$ and hence $V^B(x; k, y)$. Thus $BR(k; y)$ is increasing in $k$: the game exhibits strategic complementarity.

*Step (c): Equilibrium comparative statics.* By Proposition 1, there is a unique equilibrium threshold $\kappa_B^*(y)$ satisfying $\kappa_B^*(y) = BR(\kappa_B^*(y); y)$. We now show $\kappa_B^*(y_2) < \kappa_B^*(y_1)$ for $y_2 > y_1$. From Step (a), $BR(k; y_2) < BR(k; y_1)$ for all $k$. Suppose for contradiction that $\kappa_B^*(y_2) \geq \kappa_B^*(y_1)$. Then by strategic complementarity (Step (b)):

$$\kappa_B^*(y_2) = BR(\kappa_B^*(y_2); y_2) < BR(\kappa_B^*(y_2); y_1) \leq BR(\kappa_B^*(y_2); y_1). \tag{46}$$

Meanwhile, $\kappa_B^*(y_1) = BR(\kappa_B^*(y_1); y_1)$. Define $G(k) \equiv BR(k; y_1) - k$. We have $G(\kappa_B^*(y_1)) = 0$ and $G(\kappa_B^*(y_2)) > 0$ (since $BR(\kappa_B^*(y_2); y_1) > \kappa_B^*(y_2) \geq \kappa_B^*(y_1)$). But $V^B(x; k, y_1)$ is strictly increasing in $x$ (by Proposition 1), which means $BR(k; y_1)$ is well-defined and continuous, so $G$ is continuous. Since $G(\kappa_B^*(y_1)) = 0$ and $G(k) \to -\infty$ as $k \to +\infty$ (because $BR$ grows slower than $k$ when $V^B \to -\infty$), $G$ must have another zero above $\kappa_B^*(y_1)$, contradicting uniqueness of the equilibrium at $y_1$.

Therefore $\kappa_B^*(y_2) < \kappa_B^*(y_1)$, i.e., $\frac{d\kappa_B^*}{dy} < 0$. Since $y = \theta + \Delta$ with $\theta$ fixed, $\frac{d\kappa_B^*}{d\Delta} < 0$.

Consequently, for $\Delta \in (0, \Delta_A)$:

$$\frac{d\bar{L}}{d\Delta} = -\frac{1}{\sigma}\phi\left(\frac{\kappa_B^* - \theta}{\sigma}\right)\frac{d\kappa_B^*}{d\Delta} > 0. \quad \square \tag{47}$$

□

## B.3 Step 3: Part (ii) - Convergence and the Tacitus Trap

**Part (ii-a): Convergence.** As $\Delta \to \infty$, $\bar{x}(\Delta) = \theta + \Delta - c \to \infty$. Since private signals are $x_i \sim N(\theta, \sigma^2)$, eventually $\bar{x}(\Delta) > \kappa_S^*(\Delta)$ for all sufficiently large $\Delta$, and the economy enters Regime C. In this regime,

$$\bar{L}(\theta; \Delta) = 1 - \Phi\left(\frac{\kappa_S^*(\Delta) - \theta}{\sigma}\right). \tag{48}$$

Moreover, as $\Delta \to \infty$, essentially every investor becomes a Skeptic (fraction $\eta \to 1$). Skeptics discard the public signal and rely solely on private signals. Therefore, the equilibrium converges to a *pure private-information game* in which every investor uses only $x_i$ to form beliefs. Denote the equilibrium threshold of this limit game by $\kappa_S^\infty$, defined by

$$\kappa_S^\infty + \mathbb{E}\left[1 - \Phi\left(\frac{\kappa_S^\infty - \theta}{\sigma}\right) \,\Big|\, x_i = \kappa_S^\infty\right] - 1 = 0. \tag{49}$$

By symmetry of the standard global game without a public signal, $\kappa_S^\infty$ is independent of $y$ and hence of $\Delta$.[2] Therefore,

$$\lim_{\Delta \to \infty} \bar{L}(\theta; \Delta) = 1 - \Phi\left(\frac{\kappa_S^\infty - \theta}{\sigma}\right) \equiv \bar{L}_S(\theta). \qquad \square \tag{50}$$

**Part (ii-b): The Tacitus Trap for Large Enough Fundamental.** At $\Delta = 0$ (Regime A), $\bar{L}(\theta; 0) = 1 - \Phi\left(\frac{\kappa_B^*(0) - \theta}{\sigma}\right)$. We show $\kappa_B^*(0) < \kappa_S^\infty$ when $\theta > \kappa_S^\infty$, which implies $\bar{L}(\theta; 0) > \bar{L}_S(\theta)$.

At $\Delta = 0$, the public signal is $y = \theta$ (unbiased). The pure private-information threshold satisfies $\kappa_S^\infty + \Phi(0) - 1 = 0$, so $\kappa_S^\infty = \frac{1}{2}$. Define the fixed-point function

$$f(k) \equiv \rho\, k + (1 - \rho)\theta + \Phi\left(\frac{(1 - \rho)(\theta - k)}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right), \tag{51}$$

so that the Believer's equilibrium condition in Regime A (with $y = \theta$) is $f(\kappa_B^*) = 1$. Evaluating at $k = \kappa_S^\infty = \frac{1}{2}$ when $\theta > \frac{1}{2}$:

$$f(\tfrac{1}{2}) = \underbrace{\tfrac{\rho}{2} + (1 - \rho)\theta}_{> \frac{\rho}{2} + \frac{1-\rho}{2} = \frac{1}{2}} + \underbrace{\Phi\left(\frac{(1 - \rho)(\theta - \frac{1}{2})}{\sqrt{\sigma^2 + \tilde{\sigma}^2}}\right)}_{> \Phi(0) = \frac{1}{2}} > 1. \tag{52}$$

Since $f(k) \to -\infty$ as $k \to -\infty$, the intermediate value theorem guarantees at least one root $\kappa_B^*$ of $f(k) = 1$ with $\kappa_B^* < \frac{1}{2}$. By Proposition 1, the equilibrium is unique, so this is the only root and $\kappa_B^*(0) < \kappa_S^\infty$. Hence,

$$\bar{L}(\theta; 0) = 1 - \Phi\left(\frac{\kappa_B^*(0) - \theta}{\sigma}\right) > 1 - \Phi\left(\frac{\kappa_S^\infty - \theta}{\sigma}\right) = \bar{L}_S(\theta). \tag{53}$$

Combined with Part (ii-a), this establishes the Tacitus Trap: for $\theta > \kappa_S^\infty$, the holding mass starts above $\bar{L}_S$ at $\Delta = 0$ and converges down to $\bar{L}_S$ as $\Delta \to \infty$, so excessive manipulation eventually drives $\bar{L}$ strictly below its unmanipulated level. $\square$

## B.4 Step 4: Part (iii) - Existence of an Interior Maximum

We show that for all $\theta \in (0, 1)$, $\bar{L}$ attains an interior maximum that strictly exceeds both $\bar{L}(\theta; 0)$ and $\bar{L}_S(\theta)$.

[2] In the standard global game with uniform prior and normal signals, $\kappa_S^\infty = \frac{1}{2}$.

**$\bar{L}$ exceeds $\bar{L}(\theta;0)$.** This follows immediately from Step 2: $\bar{L}(\theta;\Delta) > \bar{L}(\theta;0)$ for $\Delta \in (0,\underline{\Delta})$.

**$\bar{L}$ exceeds $\bar{L}_S(\theta)$.** Consider $\Delta_1 \equiv \kappa_S^\infty + c - \theta > 0$, at which $\bar{x}(\Delta_1) = \kappa_S^\infty$. For $\Delta$ slightly below $\Delta_1$, we have $\bar{x}(\Delta) < \kappa_S^\infty$. If this $\Delta$ falls in Regime B (i.e., $\kappa_B^*(\Delta) < \bar{x}(\Delta) \leq \kappa_S^*(\Delta)$), then

$$\bar{L}(\theta;\Delta) = 1 - \Phi\left(\frac{\bar{x}(\Delta) - \theta}{\sigma}\right) > 1 - \Phi\left(\frac{\kappa_S^\infty - \theta}{\sigma}\right) = \bar{L}_S(\theta), \tag{54}$$

since $\bar{x}(\Delta) < \kappa_S^\infty$. If instead this $\Delta$ falls in Regime A ($\bar{x}(\Delta) \leq \kappa_B^*(\Delta)$), then by continuity of $\bar{L}$ at the Regime A/B boundary, $\bar{L}$ in late Regime A converges to $\bar{L}$ at the start of Regime B, which exceeds $\bar{L}_S$ as shown above. In either case, there exists $\Delta_0 > 0$ with $\bar{L}(\theta;\Delta_0) > \bar{L}_S(\theta)$.

**Interior maximum.** Let $L^{\max} = \sup_{\Delta \geq 0} \bar{L}(\theta;\Delta)$. By the above, $L^{\max} > \max\{\bar{L}(\theta;0), \bar{L}_S(\theta)\}$. Since $\bar{L}(\theta;\Delta) \to \bar{L}_S(\theta) < L^{\max}$ as $\Delta \to \infty$ (Step 3), there exists $M > 0$ such that $\bar{L}(\theta;\Delta) < L^{\max}$ for all $\Delta > M$. On the compact interval $[0, M]$, $\bar{L}$ is continuous and attains its supremum at some $\Delta^* \in [0, M]$. Since $\bar{L}(\theta;0) < L^{\max}$ and $\bar{L}(\theta;M) < L^{\max}$, we must have $\Delta^* \in (0, M)$ with

$$\bar{L}(\theta;\Delta^*) = L^{\max} > \max\{\bar{L}(\theta;0),\ \bar{L}_S(\theta)\}. \tag{55}$$

This completes the proof. □

# C Local Qwen Scoring: Technical Setup and Prompt Templates

## C.1 Technical Setup: Qwen2.5-14B-Instruct-AWQ

We run all textual annotation with a local large language model, **Qwen2.5-14B-Instruct-AWQ**, served by `vLLM`. The AWQ quantization enables stable and memory-efficient inference while preserving instruction-following quality.

**Inference wrapper (ChatML).** All requests are formatted in a standard ChatML wrapper:

```
<|im_start|>system
{SYSTEM_PROMPT}
<|im_end|>
<|im_start|>user
{USER_PROMPT}
<|im_end|>
<|im_start|>assistant
```

**Decoding parameters.** For *Guba* comments (private narratives), we use low-temperature decoding to reduce randomness while retaining robustness:

- temperature = 0.1, max_tokens = 1024, stop = {<|im_end|>}.

For news (public narratives), we use deterministic decoding for annotation consistency:

- temperature = 0.0, max_tokens = 512, stop = {<|im_end|>}.

## C.2 Prompt Templates

### C.2.1 Private Narratives (*Guba* Comments): System Prompt

```
你是一位专门研究中国金融市场的行为金融学家。你的任务是A股散户的社交媒体评论中提取散户情绪

你需要分析每条评论，并提取两个核心指标：
1. Sentiment (S)：投资者对未来股价方向的预期。范围 [-1.0, 1.0]。
   - -1.0：极度看空/绝望/清仓/止损/预言大跌。
   -  0.0：中性/观望/无关内容/纯粹的情绪发泄但无方向/广告。
   - +1.0：极度看多/贪婪/满仓/抄底/预言大涨。

2. Conviction (C)：投资者对自己判断的“确信度”。范围 [0.0, 1.0]。
   - 0.0 - 0.3：低确信度，提问、猜测、模棱两可。
   - 0.4 - 0.7：中确信度，普通观点或情绪描述。
   - 0.8 - 1.0：高确信度，命令式语气、断言、“立贴为证”等。

特殊词汇（2015年语境）：
-“关灯吃面”：极度悲观（-1.0），确信度中等。
-“侠之大者，为国接盘”：通常是反讽，看空（-0.8）。
-“满仓干”：极度看多（+1.0），高确信度（0.9）。
-“天台见”：绝望/极度看空（-1.0）。
-“国家队”：根据上下文判断方向。

注意：如果用户提到'主力在压盘'、'骗筹码'或'洗盘'，通常意味着用户认为后市会涨，应标记为看多

输出必须是合法 JSON。
```

### C.2.2 Private Narratives (*Guba* Comments): User Prompt Template

```
请分析以下评论列表。返回一个 JSON 对象，键为 "results"，值为列表。
列表中的每个元素需要提供："id"，"S"，"C"，"reason"（可选简短理由）。

输入数据：
{payload}

期望输出示例：
{
  "results": [
    {"id": 101, "S": -1.0, "C": 0.9, "reason": "命令式语气要求卖出"},
```

```
    {"id": 102, "S": 0.2, "C": 0.1, "reason": "疑问句，不确定"}
  ]
}
```

不要包含 Markdown 代码块标记。

### C.2.3 Public Narratives (News): System Prompt (Year-Specific)

```
你是一个客观、严谨的金融新闻数据标注员。
你精通{YEAR}年中国A股市场的历史背景及相关术语。
你的任务是阅读新闻文本，提取客观的文本属性。
不要进行过度的理论解读，只根据文本字面含义和语气和明确的言外之意进行标注。
请严格以 JSON 格式输出结果，不要包含Markdown标记。
```

### C.2.4 Public Narratives (News): User Prompt Template

```
请分析以下财经新闻片段，提取关键信息特征。

新闻来源：{source_name}
文本内容：
{text}

请输出如下 JSON 数据（不要包含Markdown格式，严禁包含 // 注释）：
{
    "sentiment_val": float,
    "certainty_score": float,
    "source_category": string,
    "topic_tag": string,
}

字段说明：
sentiment_val:
    -1.0：明确看空/利空消息
     0.0：中性/事实陈述/多空平衡
     1.0：明确看多/利好消息

certainty_score:
    0.0：传言/猜测/不确定性高（如"据传"、"或将"）
    0.5：基于分析或有条件判断
    1.0：官方正式文件/既定事实/毫无歧义

source_category（仅选其一）：
    "Official_Gov", "State_Media", "Market_Inst", "Public_Media", "Rumor_Buzz"

topic_tag（仅选其一）：
```

```
"Deleveraging", "Policy_Rescue", "Market_Action", "Macro_Economy", "Other"
```