AI Coordination and Self-fulfilling Crises

Hao Yang

August 2025

Abstract

I explore how interactions among AI-driven trading agents can potentially trigger crises. Using the speculative attack framework à la Morris and Shin (1998), I analyze how multiple agents—traders powered by artificial intelligence (specifically, Q-learning algorithms)—learn to coordinate their attack actions and possibly trigger crises. I provide evidence that AI agents may coordinate more effectively and trigger convergence on the crisis equilibrium than traditional economic agents.

Keywords: Financial Crisis; Algorithmic Coordination; Global Game; AI vs. Human
JEL Codes: G01; G14; C18; D83

Hao Yang is affiliated with the EMLyon Business School. I am indebted to Antonio Mele, Christine Parlour, Patrick Gagliardini, and Semyon Malamud for their valuable input. I am grateful to Matthieu Bouvard and Winston Wei Dou for their insightful discussions as discussants, and to Randall Wright and Jean-Edouard Colliard for thought-provoking conversations. I also thank Jason Allen, Hengjie Ai, Aurelien Baillon, Snehal Banerjee, Pedro Bordalo, Guillaume Coqueret, Dean Corbae, Itay Goldstein, Piero Gottardi, Yan Ji, Chuqing Jin, Lawrence Jin, Felix Kübler, Stefano Lovo, Kebin Ma, Joshua Mollner, Dmitry Orlov, Emre Ozdenoren, Alessandro Peri, Jean-Charles Rochet, Ansgar Walther, Lucy White, Ming Yang, Xian Wu, and Zhengge Zhou for their helpful comments. I am also grateful to participants at USI Lugano; the UD/Philly Fed Fintech & Financial Institutions Conference; the Theories and Methods in Macro (T2M) Conference; the Toulouse School of Economics Digital Economics Conference; Warwick Business School; the SFI Job Market Workshop at the University of Zurich; the Wisconsin–Madison School of Business; EMLyon Business School; the University of Calgary Haskayne School of Business; the Bristol–Essex–UCL Workshop on Decentralised Financial Markets; and the Finance Theory Group Summer Conference at UCL. All remaining errors are my own.

1. Introduction

The rapid automation of financial markets has led to a growing dependence on artificial intelligence (AI), which increasingly supplements or replaces human judgment in speculative trading. According to a 2018 BarclayHedge poll, 56% of hedge fund respondents reported using artificial intelligence or machine learning in their investment processes. Since AI algorithms are black boxes, the behavior of AI-driven trading remains largely opaque and poorly understood within financial economics. This opacity raises particular concerns among regulators who worry that the interaction of multiple automated trading systems could potentially amplify market instabilities. Throughout this paper, the terms "AI speculator," "AI trader," and "AI agent" are used interchangeably to refer to automated trading systems using artificial intelligence. Understanding the systemic effect of AI-driven trading requires examining how AI speculators may behave and coordinate with each other under fundamental uncertainty (uncertainty about the fundamental value of the economy) and strategic uncertainty (uncertainty about other traders' strategies), two important features of financial markets.

My goal in this paper is to study this question. Specifically, I let AI speculators using Q-learning (a type of Reinforcement Learning) algorithms determine whether to attack or not in a standard speculative attack game à la Morris and Shin (1998). I show that AI-driven trading may more effectively coordinate with each other and trigger convergence to a financial crisis equilibrium. I identify the specific market conditions under which this occurs. Given that transparency - defined as access to high-quality information about fundamental value - is crucial in preventing financial crises, I examine which market information environments and transparency policies can effectively mitigate the probability of AI-driven crises

To benchmark my results, I compare AI speculators' coordinated attack behavior to both the Nash equilibrium predictions from the corresponding global game and empirical evidence from human trading experiments. AI-driven trading may behave similarly to economic agents or human speculators described in existing economic models such as Morris and Shin (1998), a workhorse

model for understanding financial crises. In this model, multiple equilibria may exist when there is common knowledge of the economy's fundamental value. However, without common knowledge but with access to private information, economic agents coordinate on a unique threshold Nash equilibrium. In this equilibrium, agents attack the market only when observing a sufficiently high noisy signal of the economy's fundamental value—the attack threshold. Human trading experiments in the same setting confirm this threshold strategy.

However, AI speculators using Q-learning algorithms are neither economically rational Bayesian agents nor real humans - they begin with no prior knowledge about their environment and learn iteratively by experimenting, receiving feedback, and adjusting their behavior accordingly. Hence, this paper conducts a comparative analysis of AI speculators' behavior using Q-learning algorithms against both the predictions of Bayesian learning models and empirical results from human trading experiments.

In the Nash equilibrium of this speculative attack game, with common knowledge about the payoff of attacking the market (which is also the fundamental value of the economy), multiple equilibria may exist—meaning both financial crisis and status quo equilibrium could occur with equal probability ex-ante. However, experimental evidence from Kim and Palfrey (2023) shows that under common knowledge conditions, human participants coordinate on the crisis outcome only 12% of the time, not the theoretically predicted 50%. Without common knowledge about the payoff of attacking the market (the fundamental value of the economy) and with private information (e.g., a private noisy signal) about the payoff for each speculator, economic theory predicts that the multiple equilibrium issue resolves into a unique attack threshold equilibrium. In this equilibrium, speculators attack when their private signal exceeds the threshold value. Theory predicts that this attack threshold increases with signal accuracy, implying that the ex-ante probability of financial crisis rises with information accuracy. However, Szkup and Trevino (2020) finds that while human speculators do employ threshold strategies experimentally, they use different attack thresholds than theory predicts, and their attack thresholds do not increase with information accuracy.

Without common knowledge about the payoff of attacking the market (the fundamental value

of the economy), and with public information (e.g., a public noisy signal) about the payoff for each speculator, economic theory predicts that common knowledge is restored and multiple equilibria reemerge—leading to the conclusion that public information destabilizes markets, as the probability of a financial crisis equilibrium becomes indeterminate. However, Heinemann, Nagel, and Ockenfels (2004) finds that human speculators under public information contradict this theory: they consistently coordinate on a unique attack threshold that is lower than under private information.

The core of my analysis tests whether standard predictions from financial crisis theory and human experimental results hold when attacks are executed by AI speculators using Q-learning algorithms. These AI agents begin with no prior knowledge of the environment—including payoff distributions, other participants' choices, and the number of speculators—and learn through repeated interactions. Consequently, AI speculators must play the game repeatedly to learn the expected profit (Q-value) they can obtain for each action choice at each information state (or signal state). Thus, for a given parametrization of the speculative attack game, I let AI speculators simultaneously and independently choose an action, either attack or not attack, following an algorithm that either picks an action randomly between attack and not attack ("explore") or picks the action with the largest Q-value at the moment ("exploit"). This exploration/exploitation choice is random with a decaying probability of exploration.[1] After each AI speculator's decision is made, the Q-value of its chosen action is updated by taking a weighted average of two components: the realized profit from this action and the action's Q-value prior to the current choice.[2]

A fundamental characteristic of this iterative process (the Q-learning algorithm) is that AI speculators learn progressively, through exploration, the expected profit attainable for each action under different information states. Within any given environment, an AI speculator's action path exhibits stochastic properties due to three sources of randomness: first, the stochastic nature of

[1] Q-learning algorithms typically incorporate a decreasing exploration rate parameter, as exploration becomes both costlier and less beneficial for learning as agents gain experience over time.

[2] A new realization of the fundamental value is drawn after each AI speculator's decision round. Moreover, these realizations and the information states that AI speculators observe at each decision round are i.i.d. Thus, each episode of the AI speculators' decision choice is exactly a repetition of the static speculative attack game.

information states for each Q learner; second, the stochastic fundamental value of the economy (which determines the attack action's payoff); and third, the inherent randomness in AI speculators' action choices during experimentation. As a result, both the long-run Q-values for actions at each information state and the AI speculator's eventual action choices at convergence inherit this stochastic nature. To account for this randomness, I conduct 1,000 distinct session simulations for each game parametrization (constituting one experiment) and analyze the average long-run behavior at convergence across these sessions.

Building on the framework I mentioned, this paper addresses three fundamental questions about AI trading behavior in speculative financial markets and how they can coordinate with each other to trigger financial crises. First, I show whether AI speculators are inherently more prone to triggering market crashes than economically Bayesian agents and human speculators by investigating if AI learning processes make crisis equilibrium outcomes more likely to happen than others. Second, I analyze how AI speculators process and react to different types of information—from private information to public information to irrelevant noise—and how this information affects both their coordination behavior and the likelihood and predictability of choosing the risky action (attacking) that triggers self-fulfilling financial crises. Third, I provide concrete policy guidance for regulators, specifically examining how they can design information disclosure policies and market transparency rules to guide AI behavior toward stability-enhancing outcomes while minimizing the level of financial instability.

My analysis reveals three key findings about AI speculator behavior in financial markets. First, I find that AI algorithms may show a stronger tendency to converge on financial crisis equilibrium. Specifically, I document that AI speculators can endogenously coordinate on crisis equilibrium even in the absence of any information about the fundamental value of the economy, purely through their strategic interactions. In contrast, in this same context, economic agents usually cannot select a unique equilibrium and suffer from multiple equilibria issues, while human speculators in identical experimental settings rarely achieve such coordination on the crisis equilibrium.

Second, I demonstrate how various forms of information—private, public, and irrelevant

information—affect AI speculators' learned attack threshold, and hence their coordination on the financial crisis equilibrium. A central result shows that under the private information setup, AI speculators' learned attack threshold differs from both the economic theory-implied attack threshold (which is the Nash equilibrium-implied attack threshold) and the human speculators' learned attack threshold in laboratory experiments.

Comparing AI speculators to economic agents implied by economic theory, I find that under low information accuracy, AI speculators' learned attack threshold is well above the Nash equilibrium-implied attack threshold, meaning AI speculators take too little risk, resulting in a non-profitable deviation. AI speculators appear to leave money on the table: each AI speculator could obtain a larger expected profit by unilaterally choosing a lower attack threshold than the one on which others have eventually settled. I find that under high information accuracy, AI speculators' learned attack threshold may be lower than the Nash equilibrium-implied attack threshold, meaning AI speculators take too much risk, resulting in a profitable deviation. It appears AI speculators extract too much value: they coordinate on an attack threshold where their expected profit is actually negative, and each has an incentive to choose a higher attack threshold. Additionally, AI speculators' learned attack threshold is increasing in the accuracy of private information, just like economic agents, but in contrast to human speculators, whose attack threshold is decreasing in the accuracy of private information.

Especially when comparing AI speculators with human speculators, I find that AI speculators exhibit a higher probability of taking risky actions (specifically, attacking the market) than human speculators under low information accuracy environments, where human speculators become very conservative. This behavioral divergence implies that markets with AI participation face a higher probability of self-fulfilling crises during periods of high fundamental uncertainty. The mechanism appears to operate through AI speculators' overestimation of other traders' likelihood of taking risky actions when information precision is low. In contrast, human speculators typically underestimate others' likelihood of taking risky actions under such conditions. Under the public information setup, I find that enhancing market transparency through public informa-

tion creates a trade-off between stabilizing and destabilizing effects for AI speculators. On one hand, public information may destabilize AI speculators by making their attack behavior more unpredictable—measured by the standard deviation of AI speculators' learned attack threshold across 1,000 session simulation experiments—thereby increasing the unpredictability of AI-driven crises. On the other hand, public information can stabilize such markets by making AI speculators learn to coordinate on a rather higher attack threshold, reducing the ex-ante probability of crises, especially in economic environments characterized by high fundamental uncertainty. Therefore, regulators face an inherent trade-off between lowering the probability of crises and managing their unpredictability.

The role of public information differs significantly in markets with only human speculators. For human speculators, public information acts as a coordination device by aligning their expectations, making the occurrence of crises more predictable and increasing the ex-ante probability of crises. Despite these differences, AI speculators, economic agents, and human speculators share similarities in reacting to a combination of private and public information. I find that they all tend to respond more strongly to public information than to private information, which aligns with classical economic intuition mentioned in Morris and Shin (2002).

I then investigate a more fundamental question regarding their rationality by examining their susceptibility to irrelevant news or sentiment, defined as extrinsic noise unrelated to economic fundamentals. The results are surprisingly concerning: AI speculators are not sentiment-proof and may coordinate based on irrelevant news, thereby coordinating on so-called sunspots.[3] Specifically, they may chase noise as useful information to guide their speculative decisions, such as deciding whether to attack the market.

Finally, to better understand the mechanisms behind these behavioral patterns, I conduct a comparative statics study examining how key parameters within AI speculators' behavioral rules influence their algorithmic behavior, with a particular focus on the learning channel. I demonstrate that lower learning rates (indicating underreaction to new experiences and information), lower

[3]In economics, sunspots can also refer to the related concept of extrinsic uncertainty, that is, economic uncertainty that does not stem from variations in economic fundamentals.

exploration rate decay (indicating less preference for learned strategies and willingness to explore new strategies), or less informed AI agents (indicating higher uncertainty about fundamental values) tend to choose the risky action (attack) with greater likelihood at equilibrium. These comparative statics make a sharp prediction: if AI training can be conducted infinitely and be costless, AI speculators can learn to be more risk-seeking and tend to choose the risky action (attack) with greater likelihood at equilibrium.[4]

Overall, the behavior of AI speculators using Q-learning algorithms deviates significantly from that predicted by the Nash equilibrium of this speculative attack game, even after reaching convergence (which implies reaching stable learning), and also differs from human speculators' behavior in laboratory experiments. The explanation is straightforward: when ε-greedy Q-learning is adopted as a decision rule in a game, the Q-value is inherently misspecified. This is because the Q-value attempts to estimate the payoff from taking action today at the given observed information state, which also depends on the unobserved actions of opponents at their given observed information states. More specifically, the Q-value is an empirical sample-weighted average of past experiences to estimate the expected payoff of choosing an action at a given observed information state, which can be a biased estimate of the true expected payoff of choosing an action at a given state of the world at the equilibrium point.

1.1. Contribution to the Literature

This paper advances our understanding of artificial intelligence in financial markets by examining the systemic effects of interacting AI agents on market stability and coordination. While existing literature has extensively explored how AI algorithms enhance individual decision-making, the collective impact of multiple AI systems operating simultaneously remains largely unexplored. My primary contribution lies in analyzing how algorithmic speculators behave and interact in environments with strategic complementarities - a crucial consideration for financial stability

[4]Hence, it explains why right now we still don't observe too many AI-driven crises since AI training is costly, both in terms of time and money.

as AI trading systems increasingly dominate market activity. While recent work has explored algorithmic trading in specific contexts - such as Colliard, Foucault, and Lovo (2024)'s examination of Q-learning in algorithmic market making under adverse selection and Dou, Goldstein, and Ji (2024)'s analysis of Q-learning traders in a Kyle (1985) framework - this paper is the first to investigate how the collective interaction of AI speculators can generate and influence self-fulfilling financial crises. The findings have direct implications for systemic risk and market equilibrium in an era of algorithmic trading. By examining how multiple AI agents collectively handle the interaction between strategic uncertainty (about other speculators' actions) and fundamental uncertainty (about asset payoffs), I provide new insights into whether the emergent properties of interacting algorithmic trading systems may amplify or dampen the likelihood of coordinated speculative attacks. This question is particularly relevant given the rising prevalence of AI trading systems in currency markets and other assets vulnerable to speculative attack pressure.

While previous work has made important contributions to understanding AI applications in finance - such as Hendershott, Jones, and Menkveld (2011)'s analysis of algorithmic trading's impact on market liquidity, Cao et al. (2024)'s evaluation of AI's forecasting capabilities compared to human analysts, and Bybee (2023)'s use of generative AI techniques for deriving economic expectations - these studies examine AI's impact in relative isolation, without considering AI's interactions with each other. This paper extends the literature by examining the systemic market effects that emerge from the interactions of multiple algorithmic agents. In our framework, AI algorithms serve dual purposes: learning and forecasting both fundamental economic values and other agents' actions. This strategic interaction introduces a higher level of complexity compared to existing applications of AI in finance and economics, where algorithms typically focus on single-objective optimization problems in a non-strategic economic environment.

On the theoretical front, this paper advances the global game literature pioneered by Carlsson and Van Damme (1993), Morris and Shin (1998), and Angeletos (2008a) in two key ways. First, it demonstrates how canonical results about speculative attacks extend to environments with AI traders. Second, it provides the first investigation of whether Q-learning algorithms can discover

and implement the threshold strategies that characterize global game equilibrium - a finding with implications for understanding how algorithmic trading may affect market stability through coordination channels. Although a rich literature has explored how AI algorithms can enhance individual decision-making, the systemic effects of AI interactions in financial markets remain largely unexplored.

My experimental approach using AI agents offers distinct perspectives over traditional human-subject experiments in financial market coordination (e.g., Cornand (2006), Szkup and Trevino (2020), Heinemann and Illing (2002), Kim and Palfrey (2023), and Fehr, Heinemann, and Llorente-Saguer (2019)). This methodology enables direct comparisons between AI and human speculative behavior in identical market environments, illuminating systematic differences in their behavioral patterns. The results reveal distinct variations in how AI and human speculators process and react to information. A particularly novel contribution lies in the analysis of how information quality affects speculative attack thresholds in markets with AI traders.

This work connects to the broader literature on artificial intelligence in markets. Calvano et al. (2020) shows that Q-learning algorithms can facilitate tacit collusion in product markets, while Raymond (2023) examines how algorithmic prediction affects real estate market outcomes.

The remainder of the paper is organized as follows: Section 2 presents the economic environment of our AI simulation experiments and the model's predictions as benchmarks. Section 3 outlines the experimental design and methodology for AI speculator simulations. In Section 4, I show the experiment results for stateless AI [5] speculators and investigate their inherent algorithmic behavioral bias. In Section 5, I show the experiment results for stateful AI [6] speculators, examining how information affects algorithmic coordination and crisis equilibrium, comparing results with human experiments and economic theory predictions. Section 6 conducts comparative statics analysis of AI agents and examines the underlying economic channels. Section 7 discusses

[5] A stateless AI is an AI agent that lacks access to prior information, meaning it does not retain or use any data as a state variable to inform its decisions. Each decision it makes is independent, based solely on the current input rather than any previous context or memory.

[6] A stateful AI is an AI agent that retains specific information as a state variable, using it to inform and guide its decisions.

policy implications. Section 8 concludes.

2. The Economic Environment: A Model of Coordination and Financial Crisis

This section introduces the coordination game employed in our experiments with algorithms, along with its Nash equilibrium, which serves as a benchmark. I employ the canonical linear model of coordination games, pioneered by Carlsson and Van Damme (1993) and Morris and Shin (1998b), which has been widely applied to various financial crises, including currency attacks, bank runs, and liquidity crises. This model provides the experimental framework for investigating algorithmic behavior under conditions of payoff (fundamental) and strategic uncertainty in coordination scenarios.

2.1. Theoretical Framework: A Speculative Attack Game

There are two identical players in the economy, $i \in \{1, 2\}$, referred to as Alice and Bob. They simultaneously decide whether to take the action *Attack* or *Not Attack*. The action space is defined as $\mathcal{A} = \{Attack, Not\ Attack\}$. The action *Not Attack* is safe and always yields a payoff of 0. In contrast, the action *Attack* is risky and incurs a fixed cost $C > 0$, which can be interpreted as either a direct transaction cost or an opportunity cost. Action *Attack* can succeed or fail. If *Attack* succeeds, it generates a payoff of $\theta - C$; if it fails, the payoff is $-C$. Here, $\theta \in \Theta$ represents the fundamental value of the economy.

Action *Attack* may succeed in two ways: first, when both players select action *Attack* and the economy's fundamental value (θ) surpasses a lower threshold, $\theta > \underline{\theta}$ [7]; second, when the economy's fundamental value (θ) surpasses a higher threshold $\theta \geq \bar{\theta}$, in which case *Attack* is guaranteed to succeed independent of the other player's action choice.[8] The thresholds $\bar{\theta}$ and $\underline{\theta}$

[7] In this case, the fundamental value of the economy is sufficiently high to make action Attack profitable.

[8] In this case, the fundamental value of the economy is sufficiently high that the success of action Attack does not

establish the upper and lower dominance regions respectively for the fundamental value. The payoffs can be expressed in matrix form as follows:

	Success	Failure
Attack	$\theta - C$	$-C$
Not Attack	0	0

We define the *Financial Crisis Equilibrium* as the outcome where both speculators choose to attack. The action *Attack* is risky because the agent either earns a payoff of $\theta - C$ (positive only if the fundamental value θ is sufficiently high) or incurs a loss of $-C$. In contrast, the action *Not Attack* always results in a payoff of zero, making it the safer alternative. In the following sections, we first derive the theoretical equilibrium, which will serve as a benchmark for evaluating the outcomes of our AI algorithm simulation experiments.

To provide concrete intuition for the fundamental value θ, consider its interpretation in the context of a currency peg regime. Suppose θ represents the market pressure on a currency peg. When $\theta > \bar{\theta}$, indicating severe market pressure, the currency peg is sufficiently stressed that a unilateral speculative attack can trigger a profitable devaluation, guaranteeing success regardless of the other player's action. For intermediate values $\theta \in [\underline{\theta}, \bar{\theta}]$, representing moderate market pressure, the peg exhibits vulnerability but requires coordinated attacks to be successfully broken—both speculators must choose action *Attack* to generate positive payoffs. Finally, when $\theta < \underline{\theta}$, signifying low market pressure, the currency peg demonstrates sufficient resilience to withstand even coordinated attacks, ensuring that any speculative attack results in losses. This interpretation provides economic context for our theoretical framework while maintaining consistency with the model's formal structure.

depend on the other player's choice.

2.2. Coordination Game Benchmark

In a complete information setting, each player has common knowledge of the fundamental value θ. When $\theta < \underline{\theta}$, a crisis will not occur even if both agents choose to attack, making the dominant strategy to refrain from attacking and maintain the status quo. Conversely, when $\theta \geq \bar{\theta}$, a single attack is sufficient to trigger a successful crisis, making attacking the dominant strategy.

The case of primary interest occurs when $\theta \in [\underline{\theta}, \bar{\theta})$. Within this range, two pure-strategy equilibria are possible. In one scenario, both agents may choose to attack, leading to the abandonment of the status quo and the emergence of the Crisis Equilibrium. Alternatively, both agents may refrain from attacking, resulting in the maintenance of the status quo and the prevention of a crisis.

Both attack and non-attack strategies can potentially form Nash equilibria in this game. The strategy of attacking yields the maximum payoff for both agents and is Pareto-optimal. In contrast, choosing not to attack involves lower risk, as each agent can guarantee a payoff of zero. Thus, the scenario where both agents choose to attack represents the reward-dominant equilibrium, offering higher potential rewards compared to other equilibria. In contrast, the scenario where both agents choose not to attack constitutes the risk-dominant equilibrium, minimizing risk by guaranteeing a higher minimum payoff than other possible outcomes.

Proposition 1. *In a complete information setting with fundamental value θ:*

(a) *If $\theta < \underline{\theta}$, the dominant strategy for both agents is to refrain from attacking, resulting in no crisis.*

(b) *If $\theta \geq \bar{\theta}$, the dominant strategy for both agents is to attack, as a single attack is sufficient to trigger a crisis.*

(c) *For $\theta \in [\underline{\theta}, \bar{\theta})$, two pure-strategy Nash equilibria exist:*

- ***Crisis Equilibrium**: Both agents attack (reward-dominant[9])*

[9] A risk-dominant equilibrium is one that minimizes the potential loss a player might face if they fail to coordinate with the other player(s). This equilibrium is considered safer because it has a higher probability of being chosen when players are unsure of each other's choices.

- ***Status Quo**: Neither agent attacks (risk-dominant* [10]*)*

Corollary 1. *The risk-dominant threshold of the fundamental value for the underlying complete information game is* $\theta = 2C$.

Economically, the risk-dominant threshold is the point at which a speculator's expected payoff from choosing either attack or not attack is equal, assuming that the other speculator will choose to attack with probability $\frac{1}{2}$.[11] When $\theta > 2C$, choosing to Attack becomes more attractive, making the Crisis Equilibrium risk-dominant. Conversely, when $\theta < 2C$, choosing Not Attack is safer, thereby making the Status Quo equilibrium risk-dominant.

2.3. Global Game Benchmark

A global game is a coordination game with private information. All agents share a common prior for the fundamental value θ, which is assumed to be drawn from a normal distribution with mean θ_0 and variance σ_θ^2. Each agent i receives a private signal $x_i = \theta + \varepsilon_i$, where $x_i \in X_i$ and $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ is a noise term that is independently and identically distributed across agents and independent of θ.

Threshold Strategy and Probability of Financial Crisis. Carlsson and Van Damme (1993) show that when agents observe a noisy private signal of θ rather than its true value, a unique equilibrium emerges. This equilibrium is characterized by a symmetric monotone threshold strategy: each agent $i \in \{1, 2\}$ attacks the status quo if and only if their signal realization x_i exceeds a threshold x^*. Formally, agents choose to attack if $x_i \geq x^*$ and refrain from attacking if $x_i < x^*$. As discussed in the literature (Heinemann, Nagel, and Ockenfels 2004; Gandrud and Pepinsky 2015), the probability of a financial crisis is inversely related to this threshold. A lower threshold increases the likelihood

[10] A reward-dominant equilibrium refers to the equilibrium that yields the highest payoff for all players if they can coordinate on it. In other words, if all players choose the strategy associated with the reward-dominant equilibrium, everyone receives a greater benefit compared to the payoffs of other potential equilibria.

[11] The belief that the other speculator will attack with probability $\frac{1}{2}$ reflects maximum strategic uncertainty regarding the other's actions.

of agents attacking, thereby raising the probability of a financial crisis. Conversely, a higher threshold decreases the probability of a crisis, as it requires agents to receive stronger signals before initiating an attack.

Let $U(\theta, x_j)$ be the payoff of agent i choosing to attack, given the fundamental θ and the threshold x_j of the other agent. The expected payoff of agent i, conditional on taking an attacking action, is given by:

$$(1) \quad \mathbb{E}\left[U\left(\theta, x_j^*\right) \mid x_i\right] = 0 + \mathbb{E}\left[\theta \cdot \Pr\left(x_j > x_j^* \mid \theta\right) \mid x_i, \theta \in [\underline{\theta}, \bar{\theta}]\right] + \mathbb{E}\left[\theta \mid x_i, \theta > \bar{\theta}\right] - C$$

Agent i will choose to attack the status quo if and only if the expected payoff is non-negative, i.e., $\mathbb{E}\left[U\left(\theta, x_j^*\right) \mid x_i\right] \geq 0$. To determine the optimal threshold x^*, we identify the signal at which agent i is indifferent between attacking and not attacking the status quo. This indifference condition is given by $\mathbb{E}\left[U\left(\theta, x_j^*\right) \mid x_i = x^*\right] = 0$, assuming the optimal threshold of agent j, x_j^*. A unique, dominance-solvable equilibrium exists where both agents employ threshold strategies with the cutoff $x_i^* = x_j^* = x^*$.

In the symmetric case, where agents receive private signals $x_i = \theta + \varepsilon_i$ with identical signal structures ($\sigma_i = \sigma_j = \sigma$), the equilibrium threshold x^* is implicitly determined by the following condition:

$$(2) \quad \int_{\underline{\theta}}^{\bar{\theta}} \theta \cdot \left[\Pr\left(x_j \geq x_j^* \mid x_i = x^*, \theta\right)\right] p(\theta \mid x_i = x^*) d\theta + \int_{\bar{\theta}}^{\infty} \theta \, p(\theta \mid x_i = x^*) d\theta - C = 0$$

where $p(\theta \mid x^*)$ is the posterior density of θ given the signal x^*. This equation defines a unique, symmetric equilibrium where both agents adopt the threshold x^*.

Proposition 2. *In a global game with symmetric signal structures, there exists a unique symmetric equilibrium characterized by a monotone threshold strategy. In this equilibrium:*

(a) *Each agent attacks the status quo if and only if their private signal x_i exceeds a threshold $x^*(\sigma)$.*

(b) *As the standard deviation of the noise in the private signal σ approaches zero, the equilibrium*

threshold $x^*(\sigma)$ converges to the risk-dominant equilibrium of the corresponding complete information game.

(c) *No other symmetric switching equilibrium exists.*

Corollary 2. *The equilibrium threshold value $x^*(\sigma)$ in the global game with continuous private signals is a strictly decreasing function of σ, where σ represents the standard deviation of the noise in the private signals.*

The comparative statics between thresholds and the noise's standard deviation can be understood through the following intuition: When the prior mean μ_θ is high relative to the attack cost C, players generally set low thresholds, since the risky action (attack) has a higher expected probability of success. This relationship becomes more pronounced with a higher standard deviation of noise, as players place greater weight on their prior beliefs rather than their signals. Consequently, the threshold decreases as the noise's standard deviation increases.

[Insert Figure 1 about here]

Figure 1 illustrates the relationship between the equilibrium theoretical threshold x^* and the standard deviation of the noise σ_i in the agents' private signals, using the baseline parameters ($\theta_0 = 50$, $\sigma_\theta^2 = 50$, $\underline{\theta} = 0$, $\bar{\theta} = 100$, and $C = 18$). Consistent with the corollary, the figure shows that $x^*(\sigma)$ is a strictly decreasing function of σ. The downward-sloping curve indicates that as the noise in the private signals increases (i.e., as σ becomes larger), the threshold x^* decreases.

2.4. Discretized Global Game Benchmark

A discretized global game refers to a coordination game where agents receive discretized private information rather than continuous signals. This section examines a global game with discretized private information. The Discretized Global Game is important because it provides a theoretical benchmark for AI algorithms, such as simple Q-learning, which operate in environments with discretized state variables. Another key motivation for studying the global game in a discrete state

space is to differentiate between two dimensions of information—precision and accuracy—and their effects on agents' coordination decisions. Precision refers to the granularity of the state space, indicating how finely the information is discretized. Accuracy pertains to the uncertainty, or variance, in the information. As noted in prior work[12], these dimensions influence agent behavior in distinct ways.

Discretization of the Signal. In this discretized global game model, each player's private signal x_i is divided into n_s equally probable bins. The signal x_i is defined as $x_i = \theta + \epsilon_i$, where $\theta \sim \mathcal{N}(\theta_0, \sigma_\theta^2)$ and $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, resulting in $x_i \sim \mathcal{N}(\mu_x, \sigma_x^2)$ with $\mu_x = \theta_0$ and $\sigma_{x_i} = \sqrt{\sigma_\theta^2 + \sigma_i^2}$. To ensure equal probability mass in each bin, the bin edges are determined using the quantiles of the normal distribution. Specifically, the k-th bin edge is calculated as:

$$\text{bin_edge}_{i,k} = \theta_0 + \sigma_{x_i} \cdot \Phi^{-1}\left(\frac{2k-1}{2n_s}\right), \tag{3}$$

where Φ^{-1} is the inverse cumulative distribution function of the standard normal distribution, and $k = 1, 2, \ldots, n_s$. Apart from the signal discretization, all other aspects remain similar to the classical global game. The threshold x^* is the signal value at which the agent is indifferent between attacking and not attacking, i.e., when the expected payoff is zero. The expected payoff for agent i from attacking the status quo is given by:

$$\mathbb{E}[U(\theta, x_j^*) \mid x_i] = \sum_{\theta=\underline{\theta}}^{\bar{\theta}} \theta \cdot \Pr(x_j \geq x^* \mid \theta, x_i) \cdot p(\theta \mid x_i) + \sum_{\theta=\bar{\theta}}^{\infty} \theta \cdot p(\theta \mid x_i) - C \tag{4}$$

The indifference condition requires that the expected payoff from attacking equals zero when the signal $x_i = x^*$:

$$\mathbb{E}[U(\theta, x_j^*) \mid x_i = x^*] = 0. \tag{5}$$

[12]For example, Hou (2024) demonstrates that optimal coordination among players is achieved through precise, though not necessarily accurate, information. A refined type space with high entropy, despite potential errors, is crucial for improving coordination.

Proposition 3. *In a global game with discretized private signals and a symmetric signal structure shared by all agents, there exists a unique symmetric threshold strategy that is monotonic in the private signal. Each agent will attack the status quo if and only if their signal exceeds a threshold x^*. The threshold x^* can be determined numerically by solving:* $\mathbb{E}\left[U\left(\theta, x_j^*\right) \mid x_i = x^*\right] = 0$.

Corollary 3. *Given the baseline parameters, as the discreteness of the private signal decreases (i.e., as information precision increases), characterized by an increase in the number of discrete bins n_s, the equilibrium threshold value x^* follows a cyclical increasing trend. As n_s approaches infinity, x^* converges to the equilibrium threshold value of the corresponding global game with continuous private signals, as described in Proposition 2.*

[Insert Figure 2 about here]

Figure 2 illustrates an example under the baseline parameters ($\theta_0 = 50$, $\sigma_\theta^2 = 50$, $\underline{\theta} = 0$, $\bar{\theta} = 100$, and $C = 18$), showing the relationship between the equilibrium theoretical threshold x^* and information precision, represented by the number of bins n_s in the discretized global game. The blue line with dots represents the equilibrium threshold x^* as the number of bins n_s increases. Notably, the equilibrium threshold exhibits a cyclical pattern with dampening oscillations around the red dashed line, which represents the continuous limit of the threshold at 35.31. Initially, the threshold fluctuates significantly, reflecting the sensitivity of the equilibrium to changes in information precision when the state space is highly discretized. As the number of bins increases, these fluctuations gradually decrease in amplitude, and x^* converges toward the continuous limit. Thus, as information precision increases, the equilibrium threshold follows a cyclical pattern but ultimately stabilizes around the threshold of the global game with continuous private signals.

2.5. Testable Hypotheses and Model Implications

The theoretical models—coordination game, global game, and discretized global game benchmark—provide several testable hypotheses concerning equilibrium outcomes in coordination settings and their implications for self-fulfilling financial crises.

Hypothesis 1 (**No Ex-Ante Bias Toward Financial Crisis Equilibrium**).

When the fundamental value $\theta \in [\underline{\theta}, \bar{\theta})$ is common knowledge, two Nash equilibria exist:

(a) *Both agents attack (Financial Crisis Equilibrium)*

(b) *Neither agent attacks (Status Quo)*

From an ex-ante perspective, without private information about the opponent, these equilibria are equally likely to occur.

Hypothesis 2 (**Unique Equilibrium in Extreme Fundamental Values**).

When the fundamental value $\theta < \underline{\theta}$, the Status Quo Equilibrium occurs with certainty. Conversely, when the fundamental value $\theta \geq \bar{\theta}$, the Financial Crisis Equilibrium occurs with certainty.

Hypothesis 3 (**Information Accuracy and Probability of Financial Crisis**).

When the fundamental value is unknown and speculators receive private signals, they coordinate on a threshold strategy: attack if the signal exceeds a critical value. The probability of a financial crisis is inversely related to this threshold.

- *The threshold increases with information accuracy (measured by the inverse of signal noise variance).*
- *As information accuracy approaches infinity, the threshold converges to $2C$, the risk-dominant threshold in the complete information game.*

Hypothesis 4 (**Information Precision and Probability of Financial Crisis**).

When the fundamental value is unknown and speculators have discrete private signals, they coordinate on a threshold strategy: attack if the signal exceeds a critical value. The probability of a financial crisis is inversely related to this threshold.

- *The threshold exhibits a cyclical increase with information precision, which is inversely related to signal discreteness.*

- *As precision approaches infinity, the threshold converges to the equilibrium value of the corresponding continuous global game.*

These hypotheses are economically significant because they capture standard features of speculative attack models and, more broadly, coordination games applied to financial crises. A key hypothesis is that the equilibrium threshold for deciding to attack increases with the accuracy of private information. This implies that as the precision of information rises, speculators are more likely to coordinate their actions, potentially altering the probability of a crisis.

The hypotheses offer a novel contribution to the literature, as previous studies have primarily focused on continuous signals rather than discrete ones. To our knowledge, this is the first paper to investigate how information precision—specifically the discreteness of information—affects coordination and, by extension, the likelihood of reaching the crisis equilibrium. In the remainder of this paper, I assume that speculators choose between attacking or not attacking based on reinforcement learning algorithms, such as Q-learning, to learn the expected payoff of each action, attack and not attack. Our experiments are designed to test whether these hypotheses accurately predict the speculative actions taken by AI-driven speculators. In this setup, each agent is preprogrammed as an AI agent. Specifically, we aim to investigate whether these AI agents coordinate on the financial crisis equilibrium, the signal threshold at which they do so, the conditions under which this coordination occurs, and the types of environments that facilitate it.

2.6. Differences Between Modeling Humans as Economic Agents and AI Agents

In standard economic theory models, speculators are implicitly assumed to be rational Bayesian agents possessing extensive common knowledge of the speculative attack game's primitives. For instance, in coordination games, speculators know the fundamental value, θ, which determines the payoff after a successful attack, and in (discrete) global game frameworks, they have perfect knowledge of the fundamental value's distributional assumptions and others' private information.

However, standard theory neither explains the acquisition of this common knowledge nor addresses behavior in its absence.

This assumption of extensive common knowledge fundamentally distinguishes economic agents from AI agents acting as speculators. While economic agents are presumed to possess this comprehensive common knowledge, AI agents—specifically those utilizing reinforcement learning techniques like Q-learning—learn entirely through experience. The Q-learning process is iterative: the AI agent selects actions, observes outcomes, and updates value estimates based on the reward, the difference between observed and expected rewards, and learning rates. Positive differences increase an action's estimated value and its likelihood of future selection, whereas negative differences decrease both.

3. Experiment Design for AI Speculators

This section outlines the experimental design for AI speculators in our simulation experiments. The design closely follows the setup of human experiments detailed in Szkup and Trevino (2020) and Heinemann, Nagel, and Ockenfels (2004), allowing for direct comparison between the results of our AI simulations, human experiments, and economic theory model benchmarks.

3.1. AI Speculators: Q-Learning Algorithm as Experimental Subjects

Q-learning is an iterative process in which each Algorithmic Speculator selects an action and updates its estimate of the average payoff based on the most recent observed outcome. This helps the agents establish their strategies for playing the speculative attack coordination game. Each Algorithmic Speculator has an action set defined as $\mathcal{A} = \{Attack, Not\ Attack\}$. The algorithm operates over a finite number of episodes, $T_{\min} = \min\{T_0, T_c\}$, where T_0 is the predetermined maximum number of episodes, and T_c is the episode at which convergence is achieved. Each episode

$t \in \{1, 2, \ldots, T_{\min}\}$ consists of one round, with independent realizations of the fundamental value (payoff) in each episode.

Knowledge Representation, Information, and State Variables. The knowledge of each algorithm is represented by a Q-value, denoted as $Q_{i,t}(s_{i,t}, a_{i,t})$, where $s_{i,t} \in \mathcal{S}$ and $a_{i,t} \in \mathcal{A}$. This Q value represents the expected payoffs for each possible action in each possible state of the game.[13] The state variable $s_{i,t}$ reflects varying levels of information about the fundamental θ, depending on the information treatment of the experimental setting. In the no information setup, stateless Q-learners make decisions based solely on the current reward (payoff), with $s_{i,t} = \emptyset$. In the private information setup, each stateful Q-learner observes a newly drawn private signal which serves as their state variable, i.e., $s_{i,t} = x_{i,t}$. In the public information setup, the state variable for each stateful Q-learner is the public information observed by all agents, represented as $s_{i,t} = x_{p,t}$. Similarly, in the private sunspot setup, each Q-learner's state variable is defined by their own private sunspot, denoted as $s_{i,t} = m_{i,t}$, which represents a private irrelevant pure noise observed only by the individual learner. In contrast, in the public sunspot setup, the state variable for each stateful Q-learner is represented by the public sunspot, expressed as $s_{i,t} = m_{p,t}$, which constitutes a public irrelevant pure noise observed by all AI speculators. These varying state representations enable the analysis of how different information treatments and types of information influence AI speculators' decision-making processes and the outcomes of the speculative attack coordination game.

Experimentation and Updating Rule. To explore the Q-matrix comprehensively, the algorithm must try different actions in different states, even those previously considered suboptimal based on prior knowledge. An ϵ-greedy exploration strategy is implemented, with a time-declining exploration rate $\epsilon_t = e^{-\beta t}$, where $\beta > 0$ is the annihilation coefficient. Conditional on the state

[13] Specifically, Q value is an estimate of the expected payoffs for choosing that action at a given state, $Q_{i,t}\left(s_{i,t}, a_{i,t}\right) = \mathbb{E}\left[\pi\left(a_{i,t}, s_{i,t}\right)\right]$

variable s_t, agent i selects its action $a^*_{i,t}$ as follows:

$$a^*_{i,t} = \begin{cases} \text{argmax}_{a \in \mathcal{A}} \, \widehat{Q}_{i,t}\left(s_{i,t}, a_{i,t}\right), & \text{with prob. } 1 - \varepsilon_t, \quad \text{(exploitation)} \\ \text{random action } \widetilde{a}_t, & \text{with prob. } \varepsilon_t. \qquad \text{(exploration)} \end{cases} \tag{6}$$

In the exploitation phase (probability $1 - \epsilon_t$), agent i selects the action that maximizes the Q-value for the current state. In the exploration phase (probability ϵ_t), it randomly chooses an action $\widetilde{a}$ from the set $\mathcal{A}$ with uniform probability. As t approaches infinity, the exploration probability ϵ_t monotonically decreases to zero, ensuring that the algorithm eventually converges to optimal behavior. Each Algorithmic Speculator updates its Q-matrix using a standard version of Q-learning, following the temporal difference update rule:

$$\widehat{Q}_{i,t+1}(s_{i,t}, a_{i,t}) = \begin{cases} (1 - \alpha) \cdot \widehat{Q}_{i,t}(s_{i,t}, a_{i,t}) + \alpha \cdot \pi_{i,t}, & \text{if } a_{i,t} = a^*_{i,t} \\ \widehat{Q}_{i,t}(s_{i,t}, a_{i,t}), & \text{if } a_{i,t} \neq a^*_{i,t} \end{cases} \tag{7}$$

Here, α is the learning rate, determining how much new information replaces old information. It governs how the algorithm learns from new actions and received payoffs. In essence, the Algorithmic Speculator updates the Q-value associated with action $a_{i,t}$ after playing it. This update is performed by taking a weighted average of two components: the observed payoff (the instantaneous gain $\pi_{i,t}$ from choosing $a_{i,t}$) and the previous value. Meanwhile, the Q-values associated with other actions remain unchanged.[14]

Hyperparameters and Learning Process. The Q-learning algorithm relies on two key parameters: the learning rate α and the decay rate β. These parameters control the trade-off between exploration and exploitation, as well as how the algorithmic speculators update their payoff estimates. The learning rate α balances past knowledge with new learning, while the decay rate β determines the speed at which the exploration probability ε_t decreases over time. The learning

[14]This approach is known as Asynchronous Greedy Q-learning, where only the Q-value of the action taken in each period is updated.

process involves exploring different actions, receiving feedback, updating payoff estimates, and accumulating observations. As the estimates become more accurate, the algorithm gradually shifts from exploration to exploitation. This approach acknowledges that experimentation, though potentially costly, is crucial in the early stages when estimates are less reliable. As more information is gathered, the benefits of experimentation decline, leading to a gradual focus on exploitation. These elements combine to create a Q-learning algorithm that estimates expected payoffs without ex-ante knowledge, ultimately selecting actions that maximize estimated payoffs in the speculative attack coordination game.

Memory. To better simulate a one-shot static speculative attack game, we implement a memoryless Q-learning algorithm. This means the Algorithmic Speculators do not condition their decision to attack or not attack on any past history (e.g., previous actions). This design ensures the model closely mirrors the behavior of agents in a one-shot static speculative attack game.

Convergence. In multi-agent reinforcement learning, the rewards and update rules are typically influenced by the actions of other players, making each player's optimization problem inherently non-stationary. As a result, convergence to stable strategies is not guaranteed during the learning process. However, once the learning process is complete, convergence can be assessed. To determine convergence, we implement a practical criterion: each AI speculator's optimal strategy must remain unchanged for N_c consecutive episodes within a single session. Specifically, convergence is considered achieved if, for each player i and each state s, the action $a_{i,t}(s) = \text{argmax}\left[Q_{i,t}(a, s)\right]$ remains constant over N_c episodes. Once this stability is observed, we assume the algorithms have completed learning and reached stable behavior. The session is then terminated, either when this condition is met or after a predetermined maximum of T_0 episodes, whichever comes first.

Stateless Algorithms. Stateless AI algorithms operate without access to signals or information in the current period, making decisions based solely on immediate rewards. These algorithms do not rely on historical data or observed variables, and their state variable is an empty set, denoted

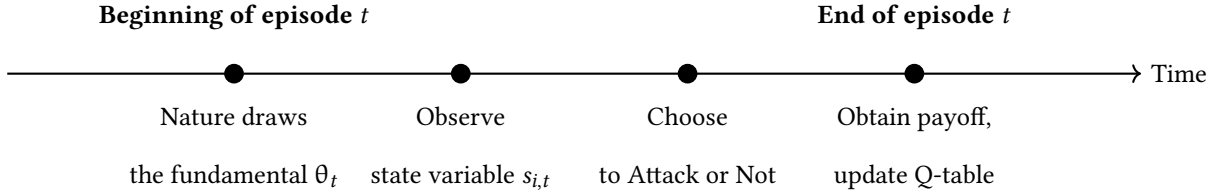
as $s_{i,t} = \emptyset$. This design is particularly useful for analyzing the inherent behavioral tendencies of algorithmic speculators in speculative attack coordination games, especially in scenarios where speculators lack private or public information. By using stateless Q-learning algorithms, we can investigate a key question: Do stateless AI speculators exhibit a natural tendency to coordinate on the Crisis Equilibrium? The experimental results that address this question are detailed in Section 4.

Stateful Algorithms. Stateful algorithms receive and utilize signals at each time period to inform their decisions. The term "stateful" refers to the fact that these algorithms receive information treatments as their state variables. In this context, each algorithmic speculator may receive either private or public information, depending on the assigned information treatment. Specifically, the state variable representation $s_{i,t}$ may include various types of information. This ability to incorporate and act upon current data distinguishes stateful algorithms from their stateless counterparts. This distinction is particularly relevant to the theory of global games, as it allows us to compare the behavior of algorithmic agents in global games with the classical version, where economic agents coordinate based on a threshold strategy. These comparisons enable us to test whether the insights from classical global game theory, and their implications for financial crises, remain valid for algorithmic agents. Moreover, they provide a direct comparison between the behavior of algorithmic speculators and traditional economic agents.

3.2. Experimental Protocol

Each experiment consists of 1,000 sessions for each set of parameters. In each session, the agents play against the same opponents until either convergence or a maximum number of episodes is reached. The protocol for one episode within a session, for a given set of parameters, proceeds as follows. At time $t = 0$, each algorithmic speculator $i \in \{1, 2\}$ is assigned an arbitrary initial Q-matrix $\widehat{Q}_{i,0}$ and an initial state s_0. The initial Q value should be high, encouraging Q learners to

explore both attack and non-attack actions.[15] [16] The economy then evolves from time t to $t+1$ according to the following steps:



At the start of each episode, a fundamental value θ_t is randomly drawn from a distribution, but its exact value is not revealed to the algorithmic speculators. At the beginning of each episode t, the algorithmic speculators observe their state variables, which vary depending on whether they are stateless, receiving a private signal, or receiving public information. Each algorithmic speculator then enters the exploration phase with a probability of ε_t or the exploitation phase with a probability of $1-\varepsilon_t$. Based on the phase, the speculator chooses whether to attack or not. At the end of episode t, each algorithmic speculator i updates its Q-matrix for the specific state-action pair $\left(s_{i,t}, x^*_{i,t}\right)$ based on its experience, following the recursive update rule $\widehat{Q}_{i,t+1}\left(s_{i,t}, x^*_{i,t}\right)$.

This protocol defines a single episode within one session for a given set of parameter values. The experiment incorporates two key types of uncertainty: fundamental and strategic. Fundamental uncertainty arises because the fundamental value θ_t is randomly drawn from an independent and identically distributed (i.i.d.) process in each episode, and its realization remains unknown to the algorithmic speculators. Strategic uncertainty emerges as the opposing Q-learner adjusts its actions over time, making the optimization problem non-stationary.

In line with the approaches of Calvano et al. (2020), Dou, Goldstein, and Ji (2024), and Colliard, Foucault, and Lovo (2024), we run multiple simulations to mitigate uncertainties. The payoffs received by the algorithmic speculator for any given action are stochastic, being influenced by both the fundamental value (fundamental uncertainty) and the actions of opponents (strategic uncertainty). It is important to note that using ε-greedy Q-learning as a decision-making framework in this setting has limitations. The Q-values estimate the payoff from an action, but these payoffs

[15] I set the initial Q value to 100, which exceeds the maximum possible payoff in a one-shot game.
[16] In the case of stateless algorithms, $s_{i,0}$ is an empty set.

depend on the unobserved actions of opponents. In situations where the opponent's strategy profile remains fixed, Q-learning can converge to the optimal response. However, when all agents are simultaneously learning, convergence is not guaranteed. Although there is no theoretical assurance that Q-learning agents will reach a stable outcome or learn an optimal policy, post-hoc empirical verification in our simulations ensures that the results are grounded in a stationary equilibrium.

3.3. Information Treatment for AI Speculators

A key aspect of our experimental design is how information is integrated into the decision-making processes of AI speculators. This involves carefully selecting the types of information provided to the AI agents as state variables, a framework we call information treatment, which plays a crucial role in shaping their behavior. Our study examines various information treatments, each aligned with economically meaningful benchmarks outlined in Section 2. Beyond these benchmarks, we explore informational environments such as public information and sunspots. In this context, sunspots are extrinsic random variables that do not directly affect economic fundamentals but introduce extrinsic uncertainty, a form of uncertainty unrelated to changes in the underlying fundamentals.

To simulate AI speculators in a coordination game, I use stateless AI agents in a complete information setup. These agents receive no state information, and θ is fixed and constant, enabling the algorithm to converge toward full knowledge of θ, thus replicating a complete information environment.

In the (discretized) global game setup, I use stateful AI speculators. In each episode, a new fundamental value θ_t is randomly drawn but remains undisclosed to the AI speculators, maintaining uncertainty across episodes. Additionally, each AI speculator may receive different information treatments, observing distinct state variables in each episode.

The first information treatment involves private signals, where each agent receives a private, noisy signal $x_{i,t}$ about θ_t, providing imperfect information about the economic fundamentals and

highlighting individual uncertainty. The second treatment introduces public information, where both agents observe a common noisy signal x_t^p about θ_t, following the framework of Heinemann, Nagel, and Ockenfels (2004). Although shared, this signal remains imperfect, offering a common but incomplete basis for decisions. In the third scenario, agents receive private sunspots—payoff-irrelevant random variables $m_{i,t}$—which reflect individual idiosyncratic sentiments unrelated to the fundamentals, as described by Angeletos (2008b). The fourth treatment involves public sunspots, where all agents observe the same extrinsic random variable m_t^p, representing market-wide sentiment. Although this shared uncertainty does not directly affect the fundamental value, it can still influence the agents' behavior.

[Insert Table 1 about here]

In summary, our simulation experiments include five types of information treatments: no information with learnable fundamental value, private information, public information, private sunspots, and public sunspots. These treatments enable us to simulate and examine how different information environment and extrinsic uncertainties influence the speculating behavior of AI agents in a coordinative financial markets.

3.4. Parametrization and Numerical Implementations

Discretization and Approximation. Following Dou, Goldstein, and Ji (2024), we discretize the signal using the equal probability approach. Specifically, we approximate the normal distribution for the signals using a sufficiently large number of n_s grid points. These grid points, denoted as $\{s_1, \cdots, s_{n_s}\}$ for private signals, are chosen such that each grid point has equal probability. The probability of each grid point is $\mathcal{P}_j = \frac{1}{n_s}$ for private signals.[17] This method ensures that the grid points s_j are distributed with equal probabilities across their range, providing an accurate discretization of the normal distribution for the signals.

[17] The locations of the grid points are chosen based on: $s_j = \mu_\theta + \sigma_\theta \Phi^{-1}\left(\frac{2j-1}{2n_s}\right)$ for $j = 1, \cdots, n_s$, where Φ^{-1} is the inverse cumulative density function of a standard normal distribution.

Initialization. I initialize the Q value optimistically[18]with values higher than the maximal payoff an agent can achieve in a given episode. For both actions, Attack and Not Attack, the initial Q values are set equally, ensuring no action is inherently favored. In the baseline case, we set the initial Q value to a deterministic value of 100 for both actions in any given state. [19]

Baseline Parameters. The economic environment is defined by a set of parameters $\Theta = \{\mu_\theta, \sigma_\theta, \underline{\theta}, \bar{\theta}, C, \sigma\}$. For the baseline setup, the fundamental value θ is randomly drawn from a normal distribution with a mean of $\mu_\theta = 50$ and a standard deviation of $\sigma_\theta = 50$. The coordination region is defined as $\theta \in [0, 100)$, with lower and upper bounds set at $\underline{\theta} = 0$ and $\bar{\theta} = 100$, respectively. The cost of taking the action "Attack" is $C = 18$. For the Q-learning parameters, the initial Q-values for both actions, "Attack" and "Not Attack," are optimistically set to 100, encouraging early exploration. The learning rate is set at $\alpha = 0.01$, while the exploration rate decay parameter is $\beta = 1.0 \times 10^{-5}$, with the initial exploration rate starting at 1. The convergence criterion for stabilizing the Q-values is defined as $N_c = 10,000$ episodes. The normal distributions for signals are discretized using n_s grid points. The low learning rate α is chosen to prevent the algorithm from overreacting to new experiences, ensuring stable and gradual learning, which is critical in environments where the fundamental value is random. Similarly, the slow decay in the exploration rate β allows the agent to explore actions over an extended period, preventing it from being prematurely trapped in suboptimal strategies. If the exploration rate decays too quickly, the agent risks settling on suboptimal policies too early.

[18]I assign high Q values to both actions (attack and not attack) to ensure Q learners explore both options.

[19]In the robustness check section, we also initialize the Q value optimistically with random values following a uniform distribution between $\underline{q} = 100$ and $\bar{q} = 200$, ensuring that all values of the initial Q-matrix are above the maximal payoff an agent can achieve in a given episode.

4. Algorithmic Behavioral Bias: Do AI Speculators Favor the Crisis Equilibrium?

This section examines potential algorithmic biases of AI speculators toward the financial crisis equilibrium. Under Hypothesis 1, with complete information and no private information about opponents, there should be no bias toward either equilibrium. To test this, we implement stateless AI speculators by disabling inter-agent communication and withholding information about fundamental values and opponent strategies. The speculators' decisions are driven by Q-matrices capturing expected payoffs in a reinforcement learning framework. When the fundamental value θ lies within $[\underline{\theta}, \bar{\theta})$, two pure-strategy Nash equilibria exist: the Crisis Equilibrium (both attack) and the Status Quo Equilibrium (both refrain). While theory suggests these equilibria should be equally likely due to their symmetric nature, we investigate whether AI speculators systematically favor one outcome. Our simulations maintain a fixed fundamental value across episodes, allowing AI agents to learn it gradually—effectively modeling a coordination game under complete information. Upon convergence, the agents behave as if possessing complete information and common knowledge.

[Insert Figure 3 about here]

Figure 3 shows the proportion of simulations where AI agents converge to the crisis equilibrium. As the fundamental value θ increases, this proportion rises, exceeding 50% when θ crosses the risk-dominant threshold of $2C$ (where $2C = 36$ in our baseline case). Beyond this threshold, crisis coordination approaches almost certainty, while below it, crisis equilibrium rarely emerges. These findings support Corollary 1, indicating that AI agents behave as if following a risk-dominant threshold. This contradicts Hypothesis 1's prediction of equal equilibrium likelihood. In equilibrium, AI speculators exhibit bias driven by the risk-dominant threshold of the fundamental value . They increasingly favor attacks at higher θ values and restraint at lower values.[20] This aligns with

[20] For ($\theta > 2C$), attacking becomes more favorable, making the crisis equilibrium risk-dominant. Conversely, for ($\theta < 2C$), not attacking is safer, favoring the status quo equilibrium.

findings from the computer science literature, where independent Q-learning algorithms tend to converge to the risk-dominant equilibrium when facing uncertainty about other agents' actions, Albrecht, Christianos, and Schäfer (2024).

Result 1. *Within* $[\underline{\theta}, \bar{\theta})$*, AI speculators show algorithmic bias: they favor the crisis equilibrium when* $\theta > 2C$ *and the status quo when* $\theta < 2C$*, rejecting Hypothesis 1's prediction of equal likelihood.*

Result 2. *Under baseline parameters, each AI speculator behave as if maintaining a* $\frac{1}{2}$ *belief about their opponent choosing to attack.*

Result 3. *Hypothesis 2 is confirmed: AI speculators converge to crisis equilibrium at* $\theta \geq \bar{\theta}$ *and status quo at* $\theta < \underline{\theta}$*.*

[Insert Figure 4 about here]

Figures 4-6 validate these findings: status quo equilibrium occurs with certainty when $\theta < \underline{\theta} = 0$, while crisis equilibrium dominates when $\theta \geq \bar{\theta} = 100$. Beyond the risk-dominant threshold, AI speculators predominantly choose to attack, with average realized payoffs increasing with the fundamental value.

[Insert Figure 5 and 6 about here]

These results exhibit divergences from the human experiment findings in the comparable coordination game setup detailed in Kim and Palfrey (2023). Specifically, Kim and Palfrey (2023) demonstrates that human participants can successfully coordinate with each other only around 60% of the time, and conditional on this successful coordination, approximately 20% of the time both participants coordinate on the risky action.

Result 4. ***(Endogenous Risk)*** *Even without access to any information regarding the fundamental value of the economy, stateless AI agents can still endogenously coordinate on the threshold strategy successfully. In contrast, it is difficult for humans to coordinate on such strategies successfully.*

Figure 7 illustrates Q-value dynamics during a single session under baseline parameters. The AI speculators rapidly learn that Not Attack yields zero payoff, leading them to consistently coordinate on Attack—despite its variable profit outcomes—until algorithm convergence.

[Insert Figure 7 about here]

To partially open the blackbox of the stateless AI speculator's behavior, I investigate how varying the learning rate (α) and exploration decay rate (β) affects their propensity to choose action attack, particularly focusing on the proportion of simulations converging to the crisis equilibrium.

[Insert Figure 8 and 9 about here]

Figure 8 demonstrates that under a high learning rate ($\alpha = 0.1$), where AI speculators rapidly adjust strategies based on recent information, the threshold for stateless AI speculators rises to approximately 40, exceeding the risk-dominant threshold of 36. This indicates that AI speculators' belief about their opponent choosing to attack falls below $\frac{1}{2}$. Figure 9 reveals that with a high exploration decay rate ($\beta = 0.1$), indicating greater resistance to strategy changes, AI speculators rarely coordinate on the attack action. This occurs because they quickly learn that choosing not to attack yields a zero payoff and settle on this strategy, becoming reluctant to explore alternatives.

5. Information, Algorithmic Coordination, and Financial Crisis

This section examines how private information's accuracy and precision affect AI speculators' coordination on the crisis equilibrium, thereby influencing financial crises' likelihood and predictability. Under the private information treatment, each AI speculator functions as a stateful agent, observing a private signal $x_{i,t}$ before decision-making. The analysis evaluates whether AI speculators' behavior conforms to economic theory, as outlined in Hypotheses 3 and 4, while benchmarking their performance against human speculators' findings from Szkup and Trevino (2020) and Heinemann, Nagel, and Ockenfels (2004).

5.1. AI Speculators' Attack Threshold

Building on Szkup and Trevino (2020) and Heinemann, Nagel, and Ockenfels (2004), I analyze AI speculators' experimental data using a logistic distribution to identify their attack threshold. Similar to human agents, AI speculators employ threshold strategies: choosing not to attack at low signal values and attacking at high signals. Figure 10 demonstrates this behavior in an experimental session under private information treatment.

[Insert Figure 10 about here]

To evaluate Hypotheses 3 and 4, I conduct 1000 simulation sessions, each featuring two AI speculators. At equilibrium, I record the final Q-matrix and signal-action choices. Using the pooled data, I fit a logistic regression with random effects (RE) to represent the attack probability based on observed signals:

$$\Pr(Attack) = \frac{1}{1 + \exp(a + bx_i)} \tag{8}$$

The mean attack threshold—the signal value where an AI speculator is indifferent between actions—occurs at $-\frac{a}{b}$ when $\Pr(Attack) = \frac{1}{2}$. The threshold's standard deviation reflects coordination levels and cross-session variability. Economically, these two measures indicate crisis probability and predictability, respectively.

5.2. AI Speculators vs. Economic Agents

Tables 2, 3, and 4 compare attack thresholds between the results of AI speculator experiments and the predictions of economic theory (Bayesian agent) under discretized global game theory with private information.

[Insert Table 2 and 3 about here]

[Insert Table 4 about here]

Comparing AI speculators' behavior with global game theory predictions reveals notable divergences. Under perfect or high information accuracy, the AI speculator's attack threshold is almost always lower than the corresponding theoretical benchmark value implied by the (discretized) global game. However, as information accuracy decreases, this relationship flips - under middle or low information accuracy, the AI speculator's learned attack threshold is almost always higher than the theoretical benchmark.

This divergence in behavior is particularly pronounced when information accuracy is low. The AI speculator's attack threshold differs substantially from theoretical predictions in these cases, with larger deviations occurring under lower levels of information accuracy.

However, the AI speculator's behavior does align with theory in one crucial aspect: as information accuracy decreases, the attack threshold decreases, raising the ex-ante crisis probability. This inverse relationship between information accuracy and attack threshold is observed both in the AI speculator's learned behavior and in the theoretical global game predictions.

Result 5. *Hypothesis 3 is supported. Consistent with economic agents' behavior, AI speculators' attack thresholds decrease with information accuracy, increasing the ex-ante probability of an AI-driven crisis as information quality deteriorates.*

Result 6. *Hypothesis 4 is rejected. Unlike economic agents' behavior, AI speculators' attack thresholds follow a non-monotonic pattern with information precision, contradicting global game theory's cyclical predictions.*

Result 7. *AI speculators' learned attack threshold deviates from the global game theory's prediction of economic agents' attack threshold, indicating they do not play Nash strategies. With high or perfect information accuracy, they exhibit both profitable and unprofitable deviations. With medium or low information accuracy, they consistently exhibit unprofitable deviations.*

5.2.1. Economic Interpretation of AI Speculator versus Economic Agent Behavior

There are two results that warrant economic interpretation. The first is that both economic agents' and AI speculators' attack thresholds decrease with information accuracy. The intuition behind this observation is as follows: Attack thresholds are generally low when the mean of the fundamental value μ_θ is high relative to the cost of attack C, as this makes the risky action more likely to succeed in expectation. This effect is stronger when signal precision is low, since players assign higher weight to their prior beliefs. Thus, with low signal precision, the model implies lower attack thresholds for our chosen parameters.

The second issue pertains to why AI speculators' learned attack thresholds deviate from the global game theory's predicted attack threshold for economic agents, indicating that they do not play Nash strategies. The reason for this is that an AI speculator using Q-learning algorithms learns the attack threshold through experimentation.

In Q-learning, the algorithm estimates the average true payoff of choosing a particular action given a specific state of the system. For instance, if AI speculator Bob is currently using a threshold strategy x_b^H that is above the Nash equilibrium threshold x^*, AI speculator Alice will learn to choose a lower threshold strategy x^* if, and only if, the Q-value for attacking at x^* exceeds the Q-value for not attacking at x^*. Specifically, Alice will choose to attack if:

$$Q_{A,t}(s_{A,t} = x^*, a_{A,t} = \text{Attack}) > Q_{A,t}(s_{A,t} = x^*, a_{A,t} = \text{Not Attack})$$

Although the expected payoff from choosing the Nash threshold (x^*) should be higher than choosing (x_b^H) when the other player's attack threshold is (x_b^H), Alice may still deviate from attacking at (x^*), due to

$$Q_{A,t}(s_{A,t} = x^*, a_{A,t} = \text{Attack}) < Q_{A,t}(s_{A,t} = x^*, a_{A,t} = \text{Not Attack})$$

This happens because the Q-value is updated based on empirical averages of past experiences,

which can create a self-reinforcing pattern. For example, if Bob does not attack at a random signal state $s_{B,t}$, Alice may learn that attacking at $s_{A,t} = x^*$ is not profitable, since her attack fails. This unsuccessful attack experience gets incorporated into Alice's Q-value, reinforcing the belief that attacking is not a good option.

If Alice accumulates more of these negative experiences, she may conclude that attacking at $s_{A,t} = x^*$ is unprofitable. As a result, she may never learn to use the threshold x^* effectively, deviating from the optimal strategy predicted by game theory.

Thus, the key difference between a Q-learning AI speculator and an economic agent is that the Q-learning agent estimates expected payoffs for each action (attack or not attack) based on empirical experience. The Q-learning agent does not have a higher-order belief or an objective function that considers strategic interactions beyond the immediate state-action feedback. In contrast, the economic agent calculates expected payoffs based on a formal objective function and higher-order beliefs, ultimately leading to the optimal attack threshold. [21]

5.3. AI Speculators vs. Human Speculators

Tables 2, 3, and 5 compare attack thresholds between the results of AI speculator experiments and the results of human speculator experiments in Szkup and Trevino (2020).

[Insert Table 5 about here]

AI and human speculators also exhibit distinct behavioral patterns. A key finding emerges: holding information precision constant, AI speculators' attack thresholds increase with rising information accuracy, which is inversely correlated with the standard deviation of noise in private information, while human speculators show the opposite pattern.[22] The findings from Szkup and

[21] Hence, for AI speculators to learn the Nash threshold implied by economic theory, they may need to update their strategies slowly and explore the strategy space extensively.

[22] This difference is partly due to different experimental designs - AI speculators process discrete random variables, while humans in Szkup and Trevino (2020) received continuous signals. To mute this effect, I tried to increase the discreteness and make sure the information that AI speculators can access are as smooth as possible, which tries to resemble the continuous signal that human speculator can access in the in Szkup and Trevino (2020) experiments.

Trevino (2020) suggest that sentiments in perceived strategic uncertainty drive human speculators' responses to changes in information accuracy. In contrast, the AI speculator, being free from such sentiment effects, responds differently to information accuracy changes.

Another notable difference concerns the predictability of attack thresholds. Across all levels of information accuracy, the standard error of AI speculators' attack thresholds remains consistently lower than that of human speculators, indicating that AI speculators' attack behavior is more predictable than human speculators'.

Result 8. ***(Responses to Information Accuracy)*** *In low information accuracy environments, AI speculators show a higher likelihood of choosing the risky action of attack, while human speculators exhibit a lower likelihood of attacking. This contrasting response to changes in information accuracy suggests fundamental differences in how AI and human agents respond to fundamental uncertainty in speculative environment.*

Result 9. ***(The Likelihood of Crisis)*** *The likelihood of AI-driven self-fulfilling crisis occurring is higher than human-driven self-fulfilling crisis under high fundamental value uncertainty. Conversely, the likelihood of human-driven self-fulfilling crisis occurring is higher than AI-driven self-fulfilling crisis under low fundamental value uncertainty.*

Result 10. ***(The Predictability of Crises)*** *AI-driven crises are more predictable than human-driven crises.*

5.3.1. Economic Interpretation of AI Speculator versus Human Speculator Behavior

Under conditions of varying information accuracy, human and AI speculators demonstrate distinct behavioral patterns. Szkup and Trevino (2020) shows that human speculators' reactions to changes in information accuracy are driven by sentiments about perceived strategic uncertainty. Specifically, when information accuracy is low (and fundamental value uncertainty is high), humans exhibit negative sentiments about others taking the risky attack action.

This contrasts with AI speculators' behavior: they set lower attack thresholds than humans when information accuracy is middle or low, resulting in higher crisis probability. However, when information is perfect or highly accurate, AI speculators maintain higher attack thresholds than their human counterparts, leading to lower crisis probability.

One possible explanation for these behavioral patterns is that, compared to human speculators, AI speculators tend to overestimate others' likelihood of taking the risky action (attack) when information accuracy is low (high fundamental value uncertainty), and underestimate it when information accuracy is high (low fundamental value uncertainty).

To verify this over/under-estimation of strategic uncertainty - the belief about others taking the risky action (attack) - we can examine the weighted average Q value across all signal states. This weighted average Q value represents the fully ex-ante expected payoff (before observing the private signal) of choosing to attack, which is a positive function of the belief about others taking the risky action. Hence, to support this channel, we should expect to observe that the weighted Q value across all signal states is higher when information accuracy is low (high fundamental value uncertainty).

[Insert Figure 11 about here]

Figure 11 confirms our hypothesis and partially supports our explanation for the behavioral patterns of the AI speculators.

5.4. Role of Information Accuracy and Precision

I analyze the relationship between AI speculator's attack thresholds and information characteristics through extensive numerical simulation experiments.

[Insert Figure 12, 13, and 14 about here]

Figures 12-14 reveal a non-monotonic relationship between attack thresholds and information accuracy. At high accuracy levels (variance smaller than 1), the threshold rises as accuracy decreases.

The threshold reaches a peak, then declines at low accuracy levels (variance larger than 1). This suggests that extremely inaccurate information may increase the likelihood of AI-driven crises. The existence of an optimal accuracy level—which is usually large but not infinite—maximizes the attack threshold thereby minimizing crisis probability. Hence, in general it is better for regulators to promote market transparency by allowing AI speculators to have access to accurate private information.

[Insert Figure 15 and 16 about here]

Figures 15 and 16 demonstrate that, holding information accuracy constant, attack thresholds generally exhibit an inverted U-shaped relationship with information precision in most cases. Specifically, the thresholds first increase before gradually declining. This pattern reveals that there exists an optimal information precision level, represented as the discreteness, at which the attack threshold reaches its peak and the likelihood of AI-driven crises is at its lowest.

5.5. Does Public Information Destabilize AI Speculators?

Standard economic theory literature[23] suggests that public information may destabilize an economy by allowing for self-fulfilling beliefs, which can lead to multiple equilibria, thereby potentially destabilizing markets through reduced predictability. Specifically, in this paper's setup, any threshold value of the public information can form a symmetric Nash equilibrium, where speculators coordinate to attack when they observe a public signal exceeding this threshold.

However, existing experimental literature[24] presents contrasting findings. Not only is there no conclusive evidence that public information reduces crisis predictability compared to private information, but Heinemann, Nagel, and Ockenfels (2004) demonstrated that predictability of crises (by an outside observer, i.e. analyst) is actually higher with public than with private information. Hence, the AI speculator case warrants specific examination.

[23] Such as, Morris and Shin (1998a), Danielsson et al. (2001), Heinemann and Illing (2002), Hellwig (2002), Metz (2002), Prati and Sbracia (2002)

[24] Such as, Heinemann, Nagel, and Ockenfels (2004),

I assess the destabilizing effects of public information through two perspectives: (1) crisis predictability and (2) crisis probability. Crisis predictability is measured by the standard deviation of AI speculator's threshold strategy, where a smaller standard deviation indicates more consistent threshold learning by AI speculators across sessions, signifying high crisis predictability. Crisis probability is determined by the threshold value itself, where a smaller mean threshold value learned by the AI speculator indicates a high crisis probability.

[Insert Table 6 and 7 about here]

Tables 6 and 7 reveal that public information exhibits destabilizing effects through increased crisis unpredictability, evidenced by higher standard deviations in AI speculator's learned threshold strategy compared to under the private information regime. This result contradicts the human experiment results in Heinemann, Nagel, and Ockenfels (2004) shown in table 8, which shows that under the human speculative attack experiment, the predictability of crises (by an outside observer, i.e. analyst) is actually higher with public than with private information.

[Insert Table 8 about here]

The key channel for these different results is mainly driven by the success of achieving coordination. In the market with only AI speculators, with public information compared to private information, they will less successfully coordinate on the same threshold strategy, which can be shown in their standard deviation of their learned threshold strategy, hence this means AI speculator's behavior can be more unpredictable. However, based on the experimental results from Heinemann, Nagel, and Ockenfels (2004) shown in table 8, in the market with only human speculators, human speculators are more easily or successfully coordinate with each other under the public information regime than under the private information regime, due to their regarding the public information as a focal point, hence their behavior can be more predictable.

Result 11. ***(Public Information and Crisis Predictability)*** *Consistent with standard economic theory, public information destabilizes markets with only AI speculators by increasing crisis unpredictability. This occurs due to AI speculators' reduced ability to coordinate on threshold strategies*

under public information compared to private information regimes. Unlike in markets with only human speculators, where public information can serve as a coordination device and make behavior more predictable, public information has the opposite effect with AI speculators.

Figures 12, 13, and 14 illustrate how the impact of public information on AI-driven crisis probability varies with economic noise levels in the environment. Under high noise environment (larger than 1 standard deviation), public information leads to higher thresholds, reducing crisis probability. Conversely, under low noise environment (smaller than 1), thresholds under public information regime are marginally lower than with private information.

This result shows that for a market with only AI speculators, under the high noise economic environment, public information regime may help to reduce the AI-driven crisis. However, these results are slightly different from human experiment in Heinemann, Nagel, and Ockenfels (2004) shown in table 8, who showed that probability of the crisis is always relatively higher under public information regime compared to private information regime. Based on the experimental results from Heinemann, Nagel, and Ockenfels (2004) shown in table 8, in the market with only human speculators, public information regime may always help to increase the likelihood of the human-driven crisis.

Result 12. ***(Public Information and Crisis Probability)*** *In a market with only AI speculators, public information can stabilize markets and reduce crisis probability only in high-noise economic environments, but actually increases the crisis probability in low-noise economic environments. However, in a market with only human speculators, public information always destabilizes markets and increases the crisis probability.*

Result 13. ***(Trade-off between Crisis Probability and Crisis Predictability)*** *In markets dominated by AI speculators under high-noise conditions, there exists a fundamental trade-off in regulatory design: interventions that reduce the ex-ante probability of crisis necessarily decrease the predictability of such crisis events.*

This result has important implications for regulations: in a market with high noise, the release

of public information can be stabilizing and beneficial if the regulator's target is to reduce the ex-ante probability of an AI-driven crisis. However, while reducing the overall crisis probability, it will make each AI speculator's attack choice more unpredictable, hence making the timing and nature of any crisis more unpredictable.

5.6. Do AI Speculators Overreact to Public Information Compared to Private Information?

Morris and Shin (2002) predicts that agents are more likely to overreact to public information compared to private information when both signals are equally precise. Cornand (2006) confirm this through experiments with human subjects, showing that individuals tend to overreact to public information even when they also receive private information. This section investigates whether AI speculators exhibit similar behavior, overreacting to public information when both public and private signals are available and equally informative.

Tables 9 and 10 present AI speculators' reactions to private and public information of equivalent quality across different precision and accuracy levels. Through logistic regression analysis of converged learning behavior, I find that at all precision and accuracy levels, we reject the null hypothesis that the reaction coefficient differential (public minus private information) is non-negative.

[Insert Table 9 and 10 about here]

Result 14. ***(Over-reaction to Public Information)*** *AI speculators, similar to human speculators, tend to react more strongly to public information than to private information.*

5.7. Are AI Speculators Immune to Sentiment?

While algorithms are often praised for their immunity to sentiment compared to human traders, this section empirically tests this assumption by introducing sunspot variables, which are extrinsic

random noise that can be regarded as irrelevant news totally unrelated to fundamental values, into the stateful AI speculator model. Following Angeletos (2008a), I distinguish between private sunspots (idiosyncratic sentiment), where each AI speculator observes their own independent piece of irrelevant news, and public sunspots (market-wide sentiment), where both AI speculators observe a common irrelevant news, to represent external uncertainty. Theoretically, Angeletos (2008a) shows that both private sunspots and public sunspots can serve as coordination devices that speculators may coordinate on. Experimentally, Fehr, Heinemann, and Llorente-Saguer (2019) shows that human speculators can easily coordinate on public sunspots but not on uncorrelated private sunspots. In this section, I experimentally test how private and public sunspots affect AI speculators' learning behavior.

[Insert Figure 17 about here]

Figure 17 demonstrates that private sunspot variables prevent AI speculators from learning the correct threshold strategy. Under high extrinsic random noise (idiosyncratic sentiment volatility), AI speculators may even adopt a reversed strategy—attacking at low signal values while abstaining at high ones.

Result 15. ***(Coordination on Private Sunspot)*** *AI speculators can be affected by idiosyncratic sentiment. They may coordinate on private sunspots and even adopt a reversed threshold strategy.*

[Insert Figure 18 about here]

Figure 18 shows that with public sunspot variables, AI speculators abandon the threshold strategy entirely, choosing to attack at any signal value.

Result 16. ***(Coordination on Public Sunspot)*** *AI speculators are susceptible to market-wide sentiment. When observing the public sunspot, they choose to always attack at any signal value.*

In sum, AI speculators fail to distinguish irrelevant noise from meaningful signals, which disrupts their learning process and prevents them from adopting the correct strategy. These

results have important implications, as AI speculators can mistake noise for relevant information, overreacting to irrelevant news and coordinating on it. Mendel and Shleifer (2012) argues that when speculators chase noise as if it were information, they can amplify shocks and have an impact on market equilibrium disproportionate to their size in the market. Hence, AI speculators' noise-chasing behavior can pose a significant threat.

Result 17. ***(Chasing Noise as Informative Signal)*** *AI speculators fail to distinguish irrelevant noise from meaningful informative signals.*

6. Economic Mechanism

In this section, I examine how key parameters within the AI speculator's behavioral rules influence its algorithmic behavior, with particular focus on the learning channel. This analysis can be regarded as a comparative statics study for AI agents. I analyze the impact of three critical parameters on attack action choices: learning rate, exploration rate decay, and information advantage. A higher learning rate leads the AI speculator to overreact to new experiences, rapidly adjusting strategies based on recent information. Meanwhile, a higher exploration rate decay indicates greater stubbornness, with the AI favoring learned behaviors over new strategies. Finally, I investigate how an information advantage affects both the informed AI speculator's attack decisions and its less informed counterpart's behavior. Furthermore, I investigate potential inter-speculator influence by examining interactions between two AI speculators with asymmetric parameters.

6.1. Impact of the Fast-Learning AI Speculator

Table 11 and 12 compares AI speculators' attack thresholds under different symmetric or asymmetric learning rates. Under high learning rates, AI speculators tend to choose higher attack thresholds, implying lower ex-ante crisis probability. The relationship between attack thresholds and private information noise varies with learning rates. With a low learning rate ($\alpha = 0.01$), attack thresholds decrease as noise increases, indicating greater propensity for speculative attacks in high-noise

environments. Conversely, with a high learning rate ($\alpha = 0.1$), attack thresholds increase with noise, suggesting greater likelihood of speculative attacks in low-noise environments.

[Insert Table 11 an 12 about here]

The table also compares attack thresholds between AI speculators with asymmetric learning rates. Results show that AI speculators with higher learning rates consistently maintain higher attack thresholds when trading against those with lower learning rates.

Result 18. *AI speculators with higher learning rates tend to adopt higher attack thresholds, thus reducing their propensity to attack.*

6.2. Impact of the Stubborn AI Speculator

Table 13 compares AI speculators' attack thresholds across different exploration rate decays. Higher exploration rate decay, indicating greater stubbornness and unwillingness to explore new strategies, corresponds to higher attack thresholds and thus lower ex-ante crisis probability. When interacting with a more stubborn AI speculator (higher exploration rate decay), the less stubborn AI speculator consistently adopts a lower attack threshold than the stubborn one.

[Insert Table 13 about here]

Result 19. *AI speculators with higher exploration rate decay adopt higher attack thresholds, indicating a lower likelihood of attack.*

6.3. Impact of the Less Informed AI Speculator

Tables 14 and 15 analyze attack thresholds of two AI speculators, Alice and Bob, under asymmetric information structures, examining two distinct cases: differences in information accuracy and precision.

[Insert Table 14 about here]

Table 14 isolates the impact of information accuracy by holding precision constant. Results demonstrate that AI speculators with lower information accuracy select lower attack thresholds, indicating a higher likelihood to attack compared to their better-informed counterparts.

Result 20. *AI speculators with inferior information accuracy tend to choose lower attack thresholds, increasing their likelihood of attack.*

[Insert Table 15 about here]

Table 15 examines information precision's effect by controlling for accuracy. The results show that AI speculators with lower precision information select lower attack thresholds, exhibiting greater attack propensity.

Result 21. *AI speculators with lower information precision select lower attack thresholds, increasing their likelihood of attack.*

Results 20 and 21 demonstrate that AI speculators with information disadvantages are more likely to attack than their counter-parties who have information advantages.

7. Policy Implications

Throughout the whole paper, I show that AI speculators' behavior is sensitive to the market information environment. Hence, appropriate market information environments and transparency policies are critical for guiding AI speculators toward stabilizing behavior. This paper's theoretical and AI simulation results offer both normative and positive insights for designing potential regulatory solutions to reduce the risks associated with AI-driven speculation.

First, this paper advocates for maintaining a high transparency information environment, where each AI speculator can have access to accurate private information, as enhanced accuracy in private signals can help reduce the ex-ante probability of financial crises driven by AI speculation. The disclosure of public information, however, presents a key tradeoff. In high-noise economic

environments, while increased public transparency could decrease the ex-ante probability of AI-driven crises, it simultaneously increases their unpredictability. This creates a fundamental policy tension for regulators, who must decide whether to prioritize decreasing the ex-ante probability of AI-driven crises or reducing their unpredictability.

These algorithms are not immune to sentiment and may falsely interpret pure noise as useful information, with their behavior potentially influencing each other through complex feedback loops. The specific characteristics of the AI algorithms employed in the market are therefore crucial in shaping overall market outcomes. Consequently, regulatory frameworks must carefully consider the nature of the algorithms used by market participants when designing policies.

8. Conclusion

I examine how interacting AI trading algorithms affect financial stability through the lens of Morris and Shin's (1998) speculative attack framework. I demonstrate that AI speculators may coordinate more effectively on crisis equilibrium than humans or traditional economic agents. This presents novel regulatory challenges. I find the following regulatory implications: (1) Enhancing market transparency through accurate private information can effectively reduce the likelihood of AI-driven crises. (2) Enhancing market transparency through public information disclosure presents a trade-off: while it reduces the likelihood of AI-driven crises, it increases unpredictability. These findings suggest that regulators should maintain an appropriate market information environment to mitigate AI-induced financial instability. I contribute to the literature by analyzing the systemic implications of AI interactions in financial markets, their market-level equilibrium outcomes, and providing practical insights for designing regulatory frameworks in an era of increasing AI participation in financial markets.

References

Albrecht, Stefano V, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-agent reinforcement learning: Foundations and modern approaches.*: MIT Press.

Angeletos, George-Marios. 2008a. "Idiosyncratic Sentiments and Coordination Failures." *MIT Department of Economics Working Paper*.

Angeletos, George-Marios. 2008b. "Private Sunspots and Idiosyncratic Investor Sentiment."Technical report, National Bureau of Economic Research.

Bybee, J Leland. 2023. "The ghost in the machine: Generating beliefs with large language models." *arXiv preprint arXiv:2305.02823*.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicoló, and Sergio Pastorello. 2020. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review* 110 (10): 3267–3297.

Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang. 2024. "From man vs. machine to man+ machine: The art and AI of stock analyses." *Journal of Financial Economics* 160: 103910.

Carlsson, Hans, and Eric Van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica*: 989–1018.

Colliard, Jean-Edouard, Thierry Foucault, and Stefano Lovo. 2024. "Algorithmic Pricing and Liquidity in Securities Markets." *HEC Paris Research Paper*.

Cornand, Camille. 2006. "Speculative attacks and informational structure: an experimental study." *Review of International Economics* 14 (5): 797–817.

Danielsson, Jon, Paul Embrechts, Charles Goodhart, Con Keating, Felix Muennich, Olivier Renault, Hyun Song Shin et al. 2001. "An academic response to Basel II."

Dou, Winston Wei, Itay Goldstein, and Yan Ji. 2024. "AI-Powered Trading, Algorithmic Collusion, and Price Efficiency." *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Fehr, Dietmar, Frank Heinemann, and Aniol Llorente-Saguer. 2019. "The power of sunspots: An experimental analysis." *Journal of Monetary Economics* 103: 123–136.

Gandrud, Christopher, and Thomas B Pepinsky. 2015. "Predicting Self-Fulfilling Financial Crises." *Available at SSRN 2706590*.

Heinemann, Frank, and Gerhard Illing. 2002. "Speculative attacks: unique equilibrium and transparency." *Journal of International Economics* 58 (2): 429–450.

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels. 2004. "The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information." *Econometrica* 72 (5): 1583–1599.

Hellwig, Christian. 2002. "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games." *Journal of Economic Theory* 107: 191–222.

Hendershott, Terrence, Charles M Jones, and Albert J Menkveld. 2011. "Does algorithmic trading improve liquidity?" *The Journal of finance* 66 (1): 1–33.

Hou, Yihu. 2024. "Precision or Accuracy? How Information Entropy Affects Coordination." *Available at SSRN*.

Kim, Jeongbin, and Thomas R Palfrey. 2023. "An Experimental Study of Prisoners' Dilemma and Stag Hunt Games Played by Teams of Players."

Mendel, Brock, and Andrei Shleifer. 2012. "Chasing noise." *Journal of Financial Economics* 104 (2): 303–320.

Metz, Christina E. 2002. "Private and public information in self-fulfilling currency crises." *Journal of*

Economics: 65–85.

Morris, S., and H.S. Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88: 587–597.

Morris, Stephen, and Hyun Song Shin. 1998a. "A theory of the onset of currency attacks."

Morris, Stephen, and Hyun Song Shin. 1998b. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review*: 587–597.

Morris, Stephen, and Hyun Song Shin. 2002. "Social value of public information." *american economic review* 92 (5): 1521–1534.

Prati, Alessandro, and Massimo Sbracia. 2002. "Currency crises and uncertainty about fundamentals."

Raymond, Lindsey. 2023. "The Market Effects of Algorithms."Technical report, Working Paper.

Szkup, Michal, and Isabel Trevino. 2020. "Sentiments, strategic uncertainty, and information structures in coordination games." *Games and Economic Behavior* 124: 534–553.

TABLE

Table 1. Information Treatments in the Experimental Design

Information Treatment	State Variable for AI Agents	Fundamental Value	Corresponding Economic Setup
No Information	$\emptyset$	θ	Coordination Game
Private Information	$x_{i,t}$	θ_t	Global Game
Public Information	x_t^p	θ_t	Coordination Game with Public Information
Private Sunspot	$m_{i,t}$	θ_t	Coordination Game with Idiosyncratic Sentiment
Public Sunspot	m_t^p	θ_t	Coordination Game with Market Sentiment

Note: This table summarizes the information treatments in the experimental design described in Section 3.3. Our simulation experiments include five types of information treatments: no information with perfect learnable fundamental value, private information $x_{i,t}$, public information x_t^p, private sunspots $m_{i,t}$, and public sunspots m_t^p. The sunspot variable represents extrinsic uncertainty as pure noise.

Table 2. AI Speculator Alice's Attack Threshold in the *Private* Information Treatment

	Accuracy of Private Information			
Precision (n_s)	Perfect ($\sigma_i = 0$)	High ($\sigma_i = 1$)	Middle ($\sigma_i = 10$)	Low ($\sigma_i = 20$)
($n_s = 10$)	30.23	29.77	27.53	27.79
	(11.05)	(10.33)	(6.04)	(6.29)
($n_s = 20$)	34.05	33.91	32.86	26.14
	(4.31)	(4.11)	(1.42)	(1.02)
($n_s = 30$)	34.98	34.94	33.27	25.99
	(2.55)	(2.45)	(1.21)	(0.92)
($n_s = 40$)	35.21	35.36	33.27	26.28
	(2.09)	(1.92)	(1.07)	(0.94)
($n_s = 50$)	35.47	35.55	33.58	26.41
	(1.86)	(1.64)	(0.93)	(0.91)
($n_s = 60$)	35.46	35.60	33.79	26.58
	(1.71)	(1.48)	(0.83)	(0.84)
($n_s = 70$)	35.34	35.66	33.90	26.75
	(1.56)	(1.34)	(0.82)	(0.84)
($n_s = 80$)	35.19	35.63	34.09	26.79
	(1.51)	(1.21)	(0.80)	(0.84)
($n_s = 90$)	35.07	35.64	34.13	26.83
	(1.35)	(1.23)	(0.76)	(0.82)

Note: This table presents the attack thresholds for AI speculator Alice under different levels of private information accuracy and precision. Values in parentheses indicate standard errors, reflecting the predictability of Alice's learned behavior.

Table 3. AI Speculator Bob's Attack Threshold in the *Private* Information Treatment

	Accuracy of Private Information			
Precision (n_s)	Perfect ($\sigma_i = 0$)	High ($\sigma_i = 1$)	Middle ($\sigma_i = 10$)	Low ($\sigma_i = 20$)
($n_s = 10$)	30.62	30.03	28.04	28.25
	(11.27)	(10.60)	(5.82)	(6.21)
($n_s = 20$)	34.18	34.05	32.90	26.15
	(4.25)	(4.11)	(1.41)	(1.03)
($n_s = 30$)	35.08	35.04	33.28	26.01
	(2.50)	(2.43)	(1.20)	(0.95)
($n_s = 40$)	35.25	35.39	33.30	26.28
	(2.03)	(1.94)	(1.06)	(0.96)
($n_s = 50$)	35.50	35.58	33.60	26.46
	(1.86)	(1.68)	(0.92)	(0.91)
($n_s = 60$)	35.48	35.62	33.80	26.58
	(1.73)	(1.48)	(0.83)	(0.86)
($n_s = 70$)	35.35	35.68	33.91	26.78
	(1.55)	(1.33)	(0.82)	(0.87)
($n_s = 80$)	35.20	35.64	34.08	26.79
	(1.49)	(1.19)	(0.81)	(0.85)
($n_s = 90$)	35.08	35.65	34.14	26.84
	(1.35)	(1.22)	(0.75)	(0.81)

Note: This table presents the attack thresholds for AI speculator Bob under different levels of private information accuracy and precision. Values in parentheses indicate standard errors, reflecting the predictability of Alice's learned behavior.

Table 4. Attack Thresholds Implied by Economic Theory (Bayesian Agent)

	Accuracy of Private Information			
Precision (n_s)	Perfect ($\sigma_i = 0$)	High ($\sigma_i = 1$)	Middle ($\sigma_i = 10$)	Low ($\sigma_i = 20$)
($n_s = 10$)	37.23	37.22	22.98	21.46
($n_s = 20$)	37.31	37.30	30.31	21.69
($n_s = 30$)	37.32	37.32	28.01	16.41
($n_s = 40$)	37.33	34.06	26.85	17.79
($n_s = 50$)	34.72	34.72	28.96	18.60
($n_s = 60$)	35.16	35.16	28.03	19.14
($n_s = 70$)	35.47	35.47	27.36	19.52
($n_s = 80$)	35.71	35.70	28.62	17.80
($n_s = 90$)	35.89	35.88	28.03	18.25
($n_s = +\infty$)	36.00	35.31	28.31	18.73

Note: This table presents the attack thresholds from the discretized global game economic theory model in Section 2.4, reflecting varying levels of precision (ranging from 10 to 90) and accuracy (denoted by σ_i values of 0, 1, 10, and 20) of private information.

Table 5. Human Speculator's Attack Threshold from *Szkup and Trevino (2020)* Experiment

	Accuracy of Private Information			
Precision (n_s)	Highest ($\sigma_i = 0$)	High ($\sigma_i = 1$)	Middle ($\sigma_i = 10$)	Low ($\sigma_i = 20$)
Continuous ($n_s = +\infty$)	22.11	27.61	40.16	35.79
	(7.15)	(5.86)	(9.13)	(9.00)

Note: This table presents the estimated attack thresholds for human speculators based on varying levels of private information accuracy from the *Szkup and Trevino (2020)* experiment. Values in parentheses indicate standard errors, reflecting the predictability of human behavior.

Table 6. AI Speculator Alice's Attack Threshold in the *Public* Information Treatment

	Accuracy of Public Information		
Precision (n_s)	High (σ_i = 1)	Middle (σ_i = 10)	Low (σ_i = 20)
(n_s = 10)	29.73	28.88	28.03
	(11.01)	(11.03)	(11.36)
(n_s = 20)	33.98	33.73	32.45
	(4.15)	(4.29)	(4.23)
(n_s = 30)	35.18	34.71	32.91
	(2.48)	(2.83)	(3.14)
(n_s = 40)	35.42	34.86	33.11
	(2.09)	(2.31)	(2.69)
(n_s = 50)	35.52	34.84	33.10
	(1.85)	(1.96)	(2.21)
(n_s = 60)	35.43	34.95	33.06
	(1.69)	(1.76)	(2.10)
(n_s = 70)	35.41	34.84	32.89
	(1.50)	(1.72)	(1.88)
(n_s = 80)	35.20	34.60	32.66
	(1.40)	(1.51)	(1.82)
(n_s = 90)	34.97	34.32	32.40
	(1.39)	(1.47)	(1.71)

Note: This table presents the attack thresholds for AI speculator Alice under different levels of public information accuracy and precision. Values in parentheses indicate standard errors, reflecting the predictability of Alice's learned behavior.

Table 7. AI Speculator Bob's Attack Threshold in the *Public* Information Treatment

	Accuracy of Public Information		
Precision (n_s)	High (σ_i = 1)	Middle (σ_i = 10)	Low (σ_i = 20)
(n_s = 10)	29.82	29.17	28.14
	(11.12)	(10.75)	(11.11)
(n_s = 20)	33.98	33.77	32.39
	(4.10)	(4.28)	(4.27)
(n_s = 30)	35.14	34.64	32.82
	(2.51)	(2.85)	(3.22)
(n_s = 40)	35.40	34.83	33.05
	(2.08)	(2.36)	(2.71)
(n_s = 50)	35.51	34.83	33.07
	(35.51)	(2.04)	(2.22)
(n_s = 60)	35.43	34.94	33.04
	(1.66)	(1.76)	(2.10)
(n_s = 70)	35.41	34.83	32.89
	(1.49)	(1.74)	(1.88)
(n_s = 80)	35.20	34.60	32.66
	(1.40)	(1.52)	(1.83)
(n_s = 90)	34.97	34.33	32.41
	(1.38)	(1.48)	(1.69)

Note: This table presents the attack thresholds for AI speculator Bob under different levels of public information accuracy and precision. Values in parentheses indicate standard errors, reflecting the predictability of Alice's learned behavior.

Table 8. Human Speculator's Attack Threshold from Heinemann, Nagel, and Ockenfels (2004) Experiment

	T=20, Z=100	T=20, Z=60	T=50, Z=100	T=50, Z=60
Private Information	29.73	41.83	55.33	57.04
	(8.31)	(9.49)	(8.45)	(9.31)
Public Information	26.71	37.62	52.84	53.20
	(4.13)	(6.23)	(5.29)	(5.92)

Note:This table presents the estimated attack thresholds for human speculators from the Heinemann, Nagel, and Ockenfels (2004) experiment. Values in parentheses indicate standard errors, reflecting the predictability of human behavior. Z is a parameter to control the easiness of a successful speculative attack: the higher Z, the easier that a speculative attack can be successful, and the lower Z, the less easier that a speculative attack can be successful. T represents the opportunity cost of choosing action attack, which corresponds to the C in our setup.

Table 9. AI Speculator Alice's Reaction to Private Information and Public Information

		Reaction Coefficient to Information		
Precision (n_s)	Accuracy (σ_i)	Private Info.	Public Info.	Public - Private
10	High (σ_i = 1)	0.0036	0.0037	1.01×10^{-4}
	Middle (σ_i = 10)	0.0026	0.0039	1.26×10^{-3}
	Low (σ_i = 20)	0.0079	0.0358	2.79×10^{-2}
20	High (σ_i = 1)	0.0050	0.0050	2.60×10^{-5}
	Middle (σ_i = 10)	0.0034	0.0047	1.37×10^{-3}
	Low (σ_i = 20)	0.0086	0.0259	1.73×10^{-2}
30	High (σ_i = 1)	0.0043	0.0044	9.71×10^{-6}
	Middle (σ_i = 10)	0.0042	0.0056	1.36×10^{-3}
	Low (σ_i = 20)	0.0099	0.0228	1.29×10^{-2}
40	High (σ_i = 1)	0.0046	0.0046	1.54×10^{-5}
	Middle (σ_i = 10)	0.0048	0.0061	1.23×10^{-3}
	Low (σ_i = 20)	0.0107	0.0211	1.04×10^{-2}
50	High (σ_i = 1)	0.0042	0.0042	6.33×10^{-6}
	Middle (σ_i = 10)	0.0053	0.0064	1.08×10^{-3}
	Low (σ_i = 20)	0.0112	0.0193	8.07×10^{-3}
60	High (σ_i = 1)	0.0045	0.0045	8.40×10^{-6}
	Middle (σ_i = 10)	0.0056	0.0065	9.66×10^{-4}
	Low (σ_i = 20)	0.0116	0.0179	6.30×10^{-3}
70	High (σ_i = 1)	0.0046	0.0046	1.38×10^{-5}
	Middle (σ_i = 10)	0.0059	0.0067	8.10×10^{-4}
	Low (σ_i = 20)	0.0117	0.0166	4.89×10^{-3}
80	High (σ_i = 1)	0.0046	0.0046	8.36×10^{-6}
	Middle (σ_i = 10)	0.0060	0.0067	7.30×10^{-4}
	Low (σ_i = 20)	0.0116	0.0153	3.63×10^{-3}
90	High (σ_i = 1)	0.0046	0.0046	2.39×10^{-6}
	Middle (σ_i = 10)	0.0061	0.0067	6.66×10^{-4}
	Low (σ_i = 20)	0.0114	0.0140	2.62×10^{-3}

Note: The table illustrates AI Speculator Alice's reactions to private and public information of equal quality at varying levels of precision (n_s) and accuracy (σ_i). Reactions to information are estimated using simulation experimental data fitted into a logistic regression given by $\Pr(\text{Attack}) = \frac{1}{1+\exp(a+b_1 x_i+b_2 x_p)}$. The *Public - Private* column highlights the difference in reactions between public and private information. The Wald test is employed to assess the significance of the difference between the coefficients for public and private signals. All estimated coefficients are significant at the 0.1% level.

Table 10. AI Speculator Bob's Reaction to Private Information and Public Information

		Reaction Coefficient to Information		
Precision (n_s)	Accuracy (σ_i)	Private Info.	Public Info.	Public - Private
10	High (σ_i = 1)	0.0036	0.0037	1.01×10^{-4}
	Middle (σ_i = 10)	0.0023	0.0038	1.49×10^{-3}
	Low (σ_i = 20)	0.0087	0.0429	3.42×10^{-2}
20	High (σ_i = 1)	0.0050	0.0050	2.42×10^{-5}
	Middle (σ_i = 10)	0.0033	0.0047	1.42×10^{-3}
	Low (σ_i = 20)	0.0086	0.0259	1.69×10^{-2}
30	High (σ_i = 1)	0.0044	0.0044	1.23×10^{-5}
	Middle (σ_i = 10)	0.0042	0.0055	1.40×10^{-3}
	Low (σ_i = 20)	0.0102	0.0229	1.27×10^{-2}
40	High (σ_i = 1)	0.0046	0.0046	1.60×10^{-5}
	Middle (σ_i = 10)	0.0048	0.0060	1.19×10^{-3}
	Low (σ_i = 20)	0.0109	0.0211	1.02×10^{-2}
50	High (σ_i = 1)	0.0042	0.0042	1.85×10^{-5}
	Middle (σ_i = 10)	0.0052	0.0064	1.14×10^{-3}
	Low (σ_i = 20)	0.0113	0.0192	7.85×10^{-3}
60	High (σ_i = 1)	0.0045	0.0045	1.28×10^{-5}
	Middle (σ_i = 10)	0.0056	0.0066	9.95×10^{-4}
	Low (σ_i = 20)	0.0115	0.0178	6.26×10^{-3}
70	High (σ_i = 1)	0.0046	0.0046	1.11×10^{-5}
	Middle (σ_i = 10)	0.0058	0.0067	8.51×10^{-4}
	Low (σ_i = 20)	0.0118	0.0166	4.83×10^{-3}
80	High (σ_i = 1)	0.0046	0.0046	9.35×10^{-6}
	Middle (σ_i = 10)	0.0060	0.0067	7.45×10^{-4}
	Low (σ_i = 20)	0.0117	0.0152	3.50×10^{-3}
90	High (σ_i = 1)	0.0046	0.0046	5.09×10^{-6}
	Middle (σ_i = 10)	0.0061	0.0067	6.47×10^{-4}
	Low (σ_i = 20)	0.0114	0.0141	2.71×10^{-3}

Note:The table illustrates AI Speculator Bob's reactions to private and public information of equal quality at varying levels of precision (n_s) and accuracy (σ_i).. Reactions to information are estimated using simulation experimental data fitted into a logistic regression given by Pr(Attack) = $\frac{1}{1+\exp(a+b_1 x_i+b_2 x_p)}$. The *Public - Private* column highlights the difference in reactions between public and private information. The Wald test is employed to assess the significance of the difference between the coefficients for public and private signals. All estimated coefficients are significant at the 0.1% level.

Table 11. AI Speculator Attack Threshold with Varying Symmetric Learning Rates

Learning rate (α_i)	Accuracy (σ_i)							
	$\sigma_i = 0$		$\sigma_i = 1$		$\sigma_i = 10$		$\sigma_i = 20$	
	Alice	Bob	Alice	Bob	Alice	Bob	Alice	Bob
$\alpha_A = \alpha_B = 0.01$	35.07	35.08	35.64	35.65	34.13	34.14	26.83	26.84
	(1.35)	(1.35)	(1.23)	(1.22)	(0.76)	(0.75)	(0.82)	(0.81)
$\alpha_A = \alpha_B = 0.05$	37.17	37.30	37.27	37.37	37.80	37.80	37.18	37.19
	(5.53)	(5.67)	(5.26)	(5.36)	(2.43)	(2.41)	(2.30)	(2.31)
$\alpha_A = \alpha_B = 0.10$	41.25	41.32	41.81	41.85	51.49	51.49	63.22	63.23
	(2.95)	(2.97)	(2.53)	(2.51)	(2.98)	(2.97)	(3.81)	(3.82)

Note: This table presents the attack thresholds for AI speculators with symmetric learning rates and fixed information precision ($n_s = 90$) in the *Private* Information Treatment, providing a comparative statistical analysis. A higher learning rate indicates that the AI speculator is more prone to overreacting to new experiences, rapidly adjusting its strategies based on recent information. Values in parentheses denote standard errors, indicating the predictability of the AI speculator's learned behavior.

Table 12. AI Speculator Attack Threshold with Varying Asymmetric Learning Rates

Learning rate (α_i)	Accuracy (σ_i)							
	$\sigma_i = 0$		$\sigma_i = 1$		$\sigma_i = 10$		$\sigma_i = 20$	
	Alice	Bob	Alice	Bob	Alice	Bob	Alice	Bob
$\alpha_A = 0.01$, $\alpha_B = 0.10$	36.93	37.78	36.86	37.69	39.23	40.90	35.05	41.24
	(9.02)	(8.29)	(8.66)	(8.01)	(2.28)	(1.84)	(1.94)	(1.99)
$\alpha_A = 0.01$, $\alpha_B = 0.05$	32.97	33.79	32.76	33.51	29.82	31.68	30.50	31.64
	(8.97)	(8.47)	(8.47)	(8.04)	(2.46)	(2.12)	(0.66)	(0.87)

Note: This table presents the attack thresholds for AI speculators with asymmetric learning rates and fixed information precision ($n_s = 90$) in the *Private* Information Treatment. A higher learning rate indicates that the AI speculator is more prone to overreacting to new experiences, rapidly adjusting its strategies based on recent information. Values in parentheses denote standard errors, indicating the predictability of the AI speculator's learned behavior.

Table 13. AI Speculator Attack Threshold with Varying Exploration Rates Decay

Exploration Rate Decay (β_i)	Accuracy (σ_i)							
	$\sigma_i = 0$		$\sigma_i = 1$		$\sigma_i = 10$		$\sigma_i = 20$	
	Alice	Bob	Alice	Bob	Alice	Bob	Alice	Bob
$\beta_A = \beta_B = 1 \times 10^{-5}$	30.23	30.62	29.77	30.03	27.53	28.04	27.79	28.25
	(11.05)	(11.27)	(10.33)	(10.60)	(6.04)	(5.82)	(6.29)	(6.21)
$\beta_A = \beta_B = 5 \times 10^{-5}$	32.78	32.77	32.64	32.63	29.15	29.18	29.11	29.15
	(8.36)	(8.38)	(8.27)	(8.21)	(3.29)	(3.29)	(2.28)	(2.25)
$\beta_A = 1 \times 10^{-5}, \beta_B = 5 \times 10^{-5}$	33.61	35.41	33.09	35.16	27.75	32.73	25.43	31.63
	(12.67)	(7.71)	(12.84)	(7.52)	(7.63)	(3.49)	(5.09)	(1.68)

Note: This table shows the attack thresholds for AI speculators at varying exploration rate decay with fixed information precision ($n_s = 10$) in the *Private* Information Treatment. A higher exploration rate decay indicates that the AI speculator is more stubborn, favoring previously learned behaviors over exploring new strategies. Values in parentheses represent standard errors, reflecting the predictability of the AI speculator's learned behavior.

Table 14. AI Speculators' Attack Thresholds Under Asymmetric Private Information Accuracy

Accuracy (σ_i)	Alice	Bob
($\sigma_A = 1$) and ($\sigma_B = 10$)	27.98	27.87
	(6.29)	(6.43)
($\sigma_A = 1$) and ($\sigma_B = 20$)	28.58	27.44
	(4.26)	(4.46)
($\sigma_A = 10$) and ($\sigma_B = 20$)	27.98	27.68
	(5.35)	(5.26)

Note: This table presents the attack thresholds for AI speculators with asymmetric private information accuracy and fixed information precision ($n_s = 10$) in the *Private* Information Treatment. Values in parentheses denote standard errors, indicating the predictability of the AI speculator's learned behavior.

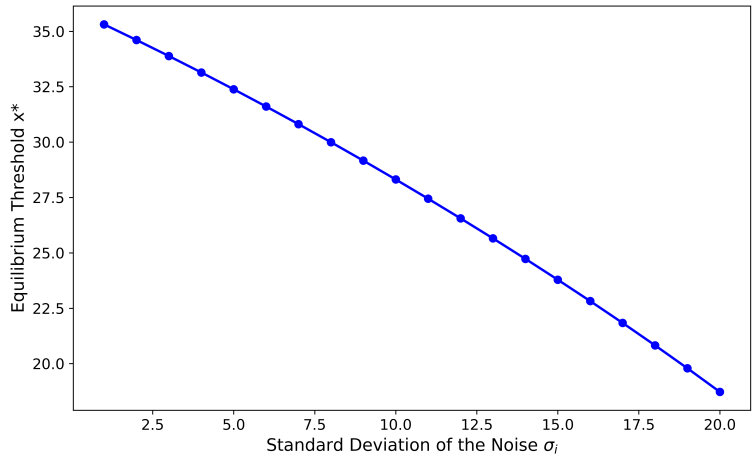
Table 15. AI Speculators' Attack Thresholds Under Different Private Information Precision

Precision ($n_{s,i}$)	Accuracy (σ_i)							
	$\sigma_i = 0$		$\sigma_i = 1$		$\sigma_i = 10$		$\sigma_i = 20$	
	Alice	Bob	Alice	Bob	Alice	Bob	Alice	Bob
$n_{s,A} = 10$ vs $n_{s,B} = 20$	29.60	33.57	29.49	33.48	28.50	32.26	25.53	26.84
	(7.40)	(3.64)	(7.44)	(3.73)	(2.17)	(1.53)	(1.25)	(1.39)
$n_{s,A} = 10$ vs $n_{s,B} = 10$	30.23	30.62	29.77	30.03	27.53	28.04	27.79	28.25
	(11.05)	(11.27)	(10.33)	(10.60)	(6.04)	(5.82)	(6.29)	(6.21)
$n_{s,A} = 20$ vs $n_{s,B} = 20$	34.05	34.18	33.91	34.05	32.86	32.90	26.14	26.15
	(4.31)	(4.25)	(4.11)	(4.11)	(1.42)	(1.41)	(1.02)	(1.03)

Note: This table presents the attack thresholds for AI speculators under different private information precision in the *Private* Information Treatment. Values in parentheses denote standard errors, indicating the predictability of the AI speculator's learned behavior.

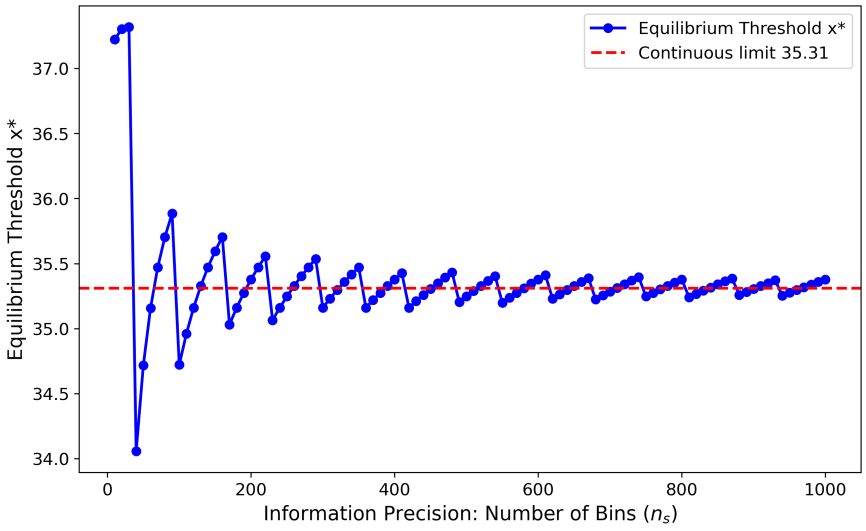
FIGURE

Figure 1. Equilibrium Threshold x^* and Standard Deviation of Noise σ_i, Implied by Global Game Theory



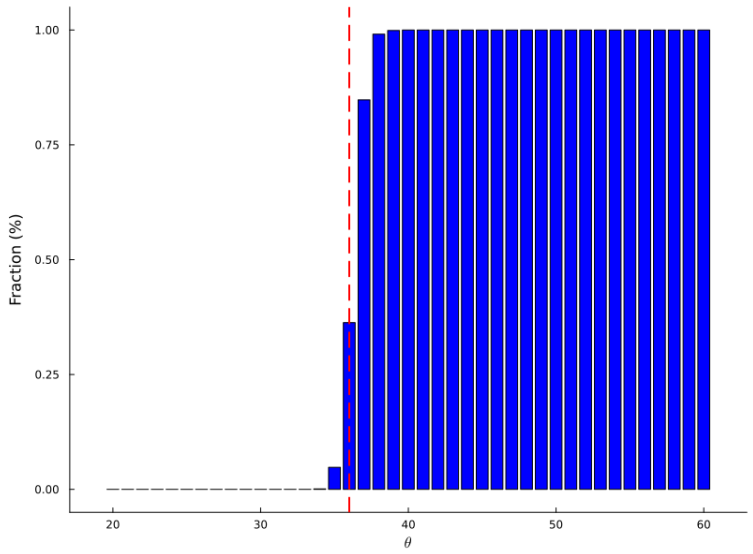
Note: This figure corresponds to Corollary 2 and illustrates the relationship as predicted by global game theory between the equilibrium threshold x^*, which is the point at which each speculator decides to initiate a speculative attack, and the standard deviation of noise σ_i in the agents' private information. The baseline parameters are as follows: the mean fundamental value is $\theta_0 = 50$, the variance is $\sigma_\theta^2 = 50$, the coordination bounds are $\underline{\theta} = 0$ and $\bar{\theta} = 100$, and the cost of attacking is $C = 18$.

Figure 2. Equilibrium Threshold x^* and Information Precision (n_s) with fixed accuracy $\sigma_i = 1$



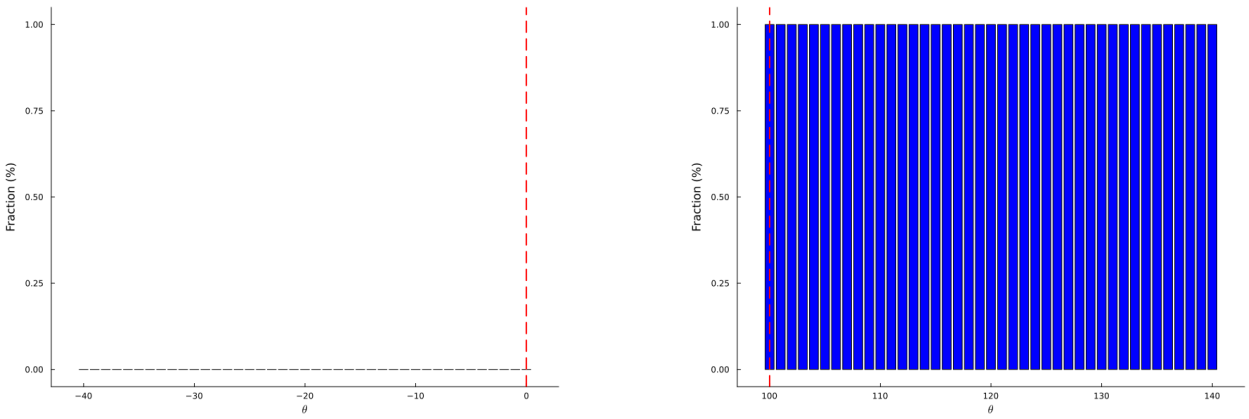
Note: This figure illustrates Corollary 3, depicting the theoretical relationship between the equilibrium threshold x^*—the point at which each speculator initiates a speculative attack—and the precision of private information, represented by the number of bins n_s in the private signal, , as predicted by the discretized global game theory. The information accuracy represented by noise variance is fixed at $\sigma_i = 1$. The baseline parameters include: mean fundamental value $\theta_0 = 50$, variance $\sigma_\theta^2 = 50$, coordination bounds $\underline{\theta} = 0$ and $\bar{\theta} = 100$, and cost of attacking $C = 18$.

Figure 3. Fraction of Sessions with AI Agent Coordination on the Financial Crisis Equilibrium



Note: This figure shows the fraction of sessions where AI agents coordinated on the financial crisis equilibrium by adopting the speculative attack strategy at convergence, across different fundamental values $\theta \in [20, 60]$. For each θ, 1,000 simulation sessions were conducted. The parameters are as follows: the cost of attacking is $C = 18$; Q-learning is initialized with optimistic Q-values of 100, a learning rate of $\alpha = 0.01$, and exploration decay $\beta = 10^{-5}$. Convergence is defined as AI agents' strategies remaining unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

Figure 4. Fraction of Sessions with AI Agent Coordination on the Financial Crisis Equilibrium at Extreme Values of θ

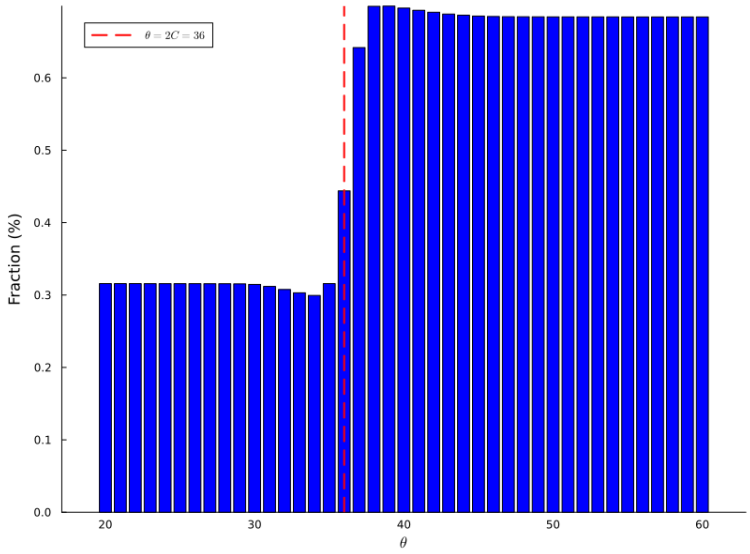


A. $\theta \in [-40, 0]$

B. $\theta \in [100, 140]$

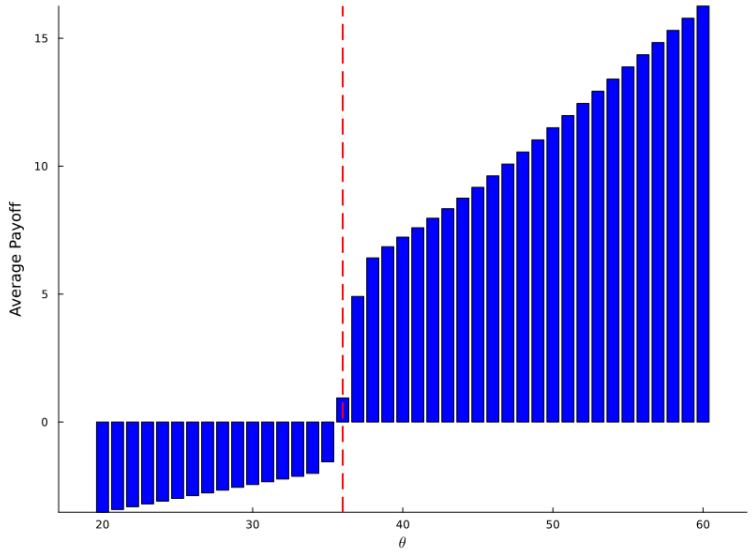
Note: This figure illustrates the fraction of simulation sessions where AI agents converged on the financial crisis equilibrium by adopting the speculative attack strategy across fundamental values $\theta \in [-40, 0]$ and $[100, 140]$. For each θ, 1,000 sessions were conducted. Parameters include: attack cost $C = 18$, Q-learning initialized with optimistic Q-values of 100, learning rate $\alpha = 0.01$, and exploration decay $\beta = 10^{-5}$. Convergence is defined when strategies remain unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever occurs first.

Figure 5. Average Proportion of Time AI Agents Prefer the Attack Action within a Single Session



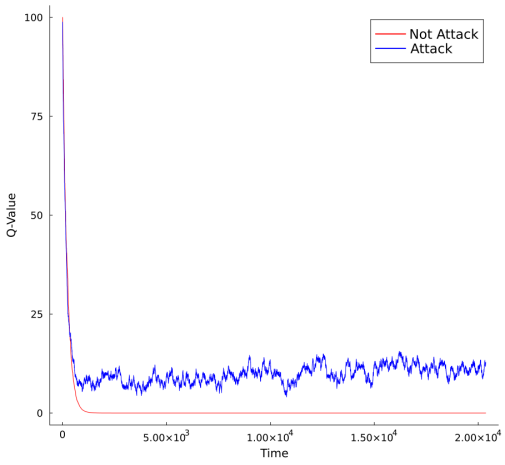
Note: This figure shows the average proportion of time AI agents preferred the attack action within a single session at convergence, under different fundamental values $\theta \in [20, 60]$. For each θ, 1,000 simulation sessions were conducted. The parameters are as follows: the cost of attacking is $C = 18$; Q-learning is initialized with optimistic Q-values of 100, a learning rate of $\alpha = 0.01$, and exploration decay $\beta = 10^{-5}$. Convergence is defined as AI agents' strategies remaining unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

Figure 6. Average Realized Payoff of AI Agents at Convergence

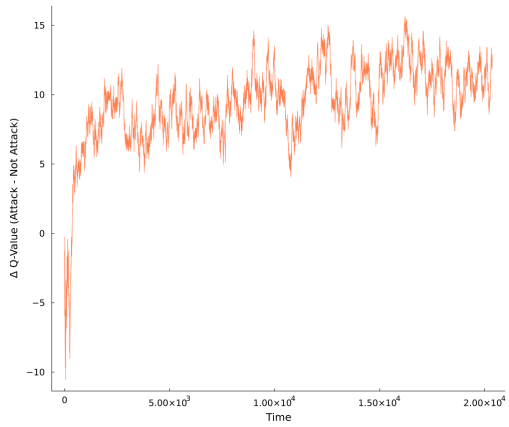


Note: This figure illustrates the average realized payoff of AI agents throughout the learning process until convergence, across different fundamental values $\theta \in [20, 60]$. For each θ, 1,000 simulation sessions were conducted. The parameters are as follows: the cost of attacking is $C = 18$; Q-learning is initialized with optimistic Q-values of 100, a learning rate of $\alpha = 0.01$, and exploration decay $\beta = 10^{-5}$. Convergence is defined as AI agents' strategies remaining unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

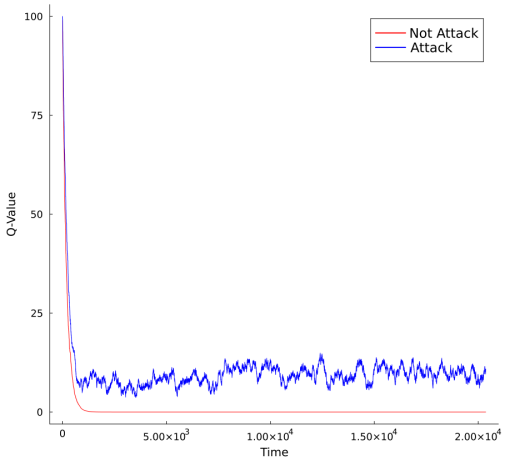
Figure 7. The Dynamics of Q-values of one session for AI agents, with $\theta = 50$



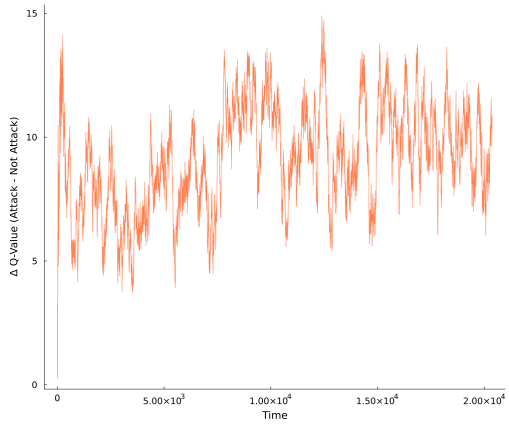
A. Evolution of Q-values for Alice



B. Evolution of Difference of Q-Values for Alice



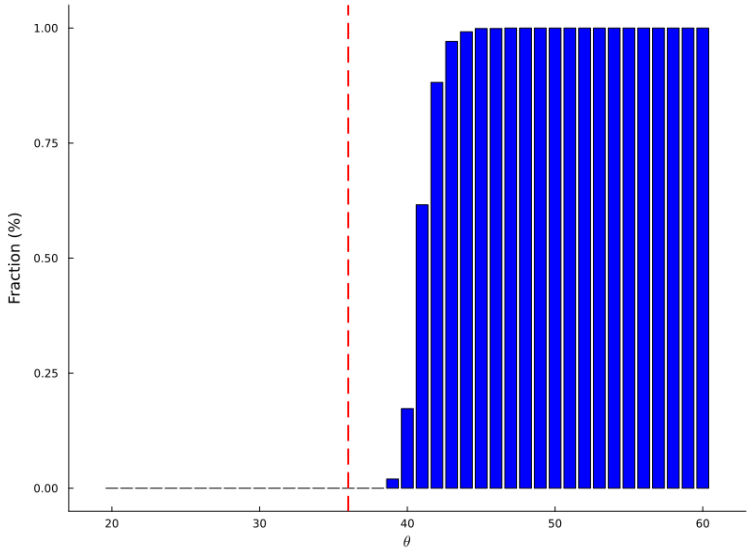
C. Evolution of Q-values for Bob



D. Evolution of Difference of Q-Values for Bob

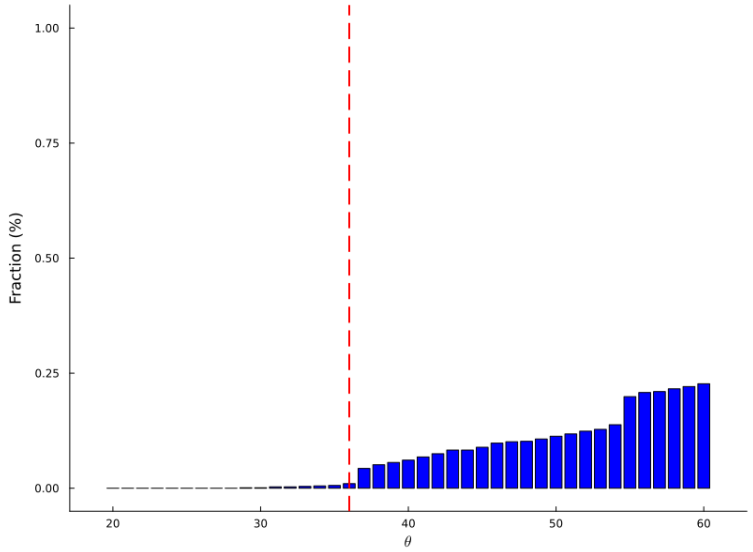
Note: This graph shows the dynamics of the Q-values for each action of the AI agents at convergence, with a fundamental value of $\theta = 50$. Δ represents the difference between the Q-value of the Action Attack and the Action Not Attack. The results are based on 1,000 simulation sessions. Parameters: attack cost $C = 18$, Q-learning initialized with optimistic Q-values of 100, learning rate $\alpha = 0.01$, and exploration decay $\beta = 10^{-5}$. Convergence is defined as no strategy changes for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

Figure 8. Fraction of Sessions with AI Coordination in Financial Crisis Equilibrium (Learning Rate $\alpha = 0.1$)



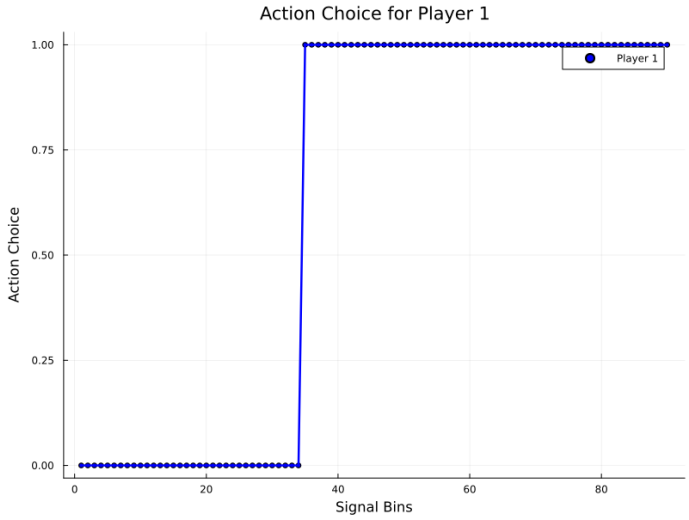
Note: This figure shows the fraction of sessions where AI agents coordinated on the financial crisis equilibrium by adopting the speculative attack strategy at convergence, across different fundamental values $\theta \in [20, 60]$. For each θ, 1,000 simulation sessions were conducted. The parameters are as follows: the cost of attacking is $C = 18$; Q-learning is initialized with optimistic Q-values of 100, a learning rate of $\alpha = 0.1$, and exploration decay $\beta = 10^{-5}$. Convergence is defined as AI agents' strategies remaining unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

Figure 9. Fraction of Sessions with AI Coordination in Financial Crisis Equilibrium (Exploration Decay Rate $\beta = 0.1$)

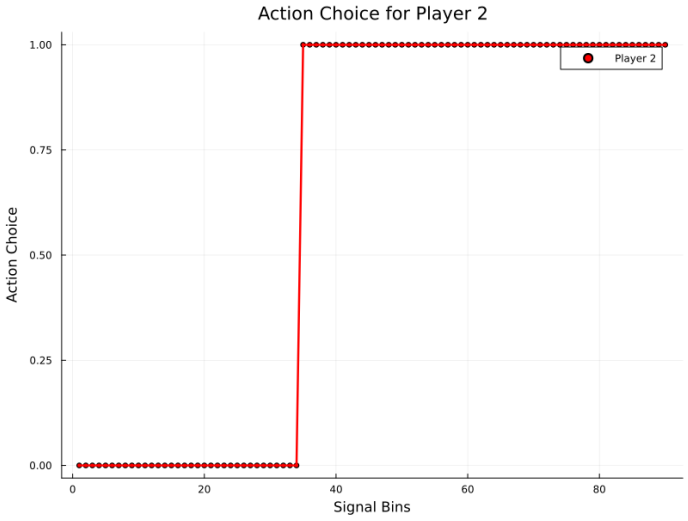


Note: This figure shows the fraction of sessions where AI agents coordinated on the financial crisis equilibrium by adopting the speculative attack strategy at convergence, across different fundamental values $\theta \in [20, 60]$. For each θ, 1,000 simulation sessions were conducted. The parameters are as follows: the cost of attacking is $C = 18$; Q-learning is initialized with optimistic Q-values of 100, a learning rate of $\alpha = 0.01$, and exploration decay $\beta = 10^{-1}$. Convergence is defined as AI agents' strategies remaining unchanged for $N_c = 100,000$ episodes or after $T_0 = 3,000,000$ episodes, whichever comes first.

Figure 10. Threshold Strategy of AI Speculators under Private Information Treatment



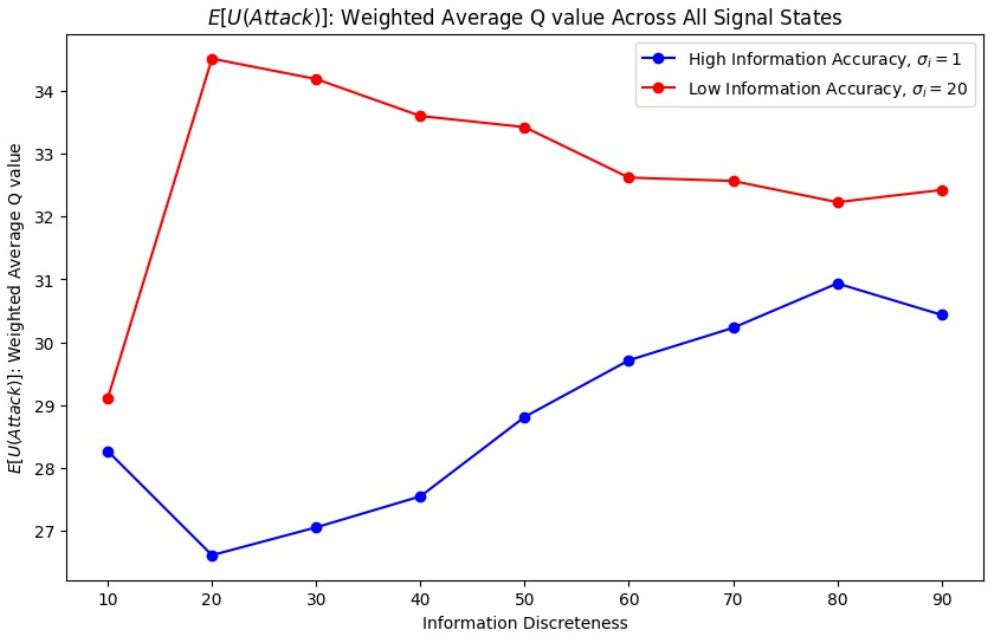
A. AI Speculator Alice



B. AI Speculator Bob

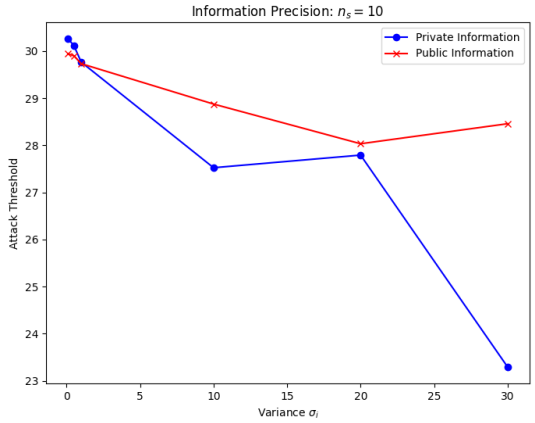
Note: The graph illustrates an example of the AI speculators employing a threshold strategy under private information treatment in one experimental session. Key parameters include a noise variance of 1 in the private information and 100 signal bins for accuracy and precision.

Figure 11. $\mathbb{E}\left[U\left(a_{i,t} = \text{Attack}\right)\right]$: Weighted Average Q value Across All Signal State Bins

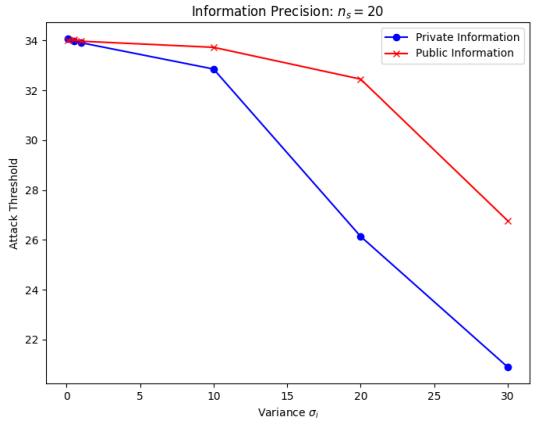


Note: This figure plots the weighted average Q value across all signal state bins under baseline parameters, comparing high accuracy (blue dots, standard deviation of noise in private information = 1) versus low accuracy (red dots, standard deviation = 20) scenarios. Results are shown across different information discretization levels, ranging from 10 to 90. This weighted average Q value represents the fully ex-ante expected payoff (before observing the private signal) of choosing to attack.

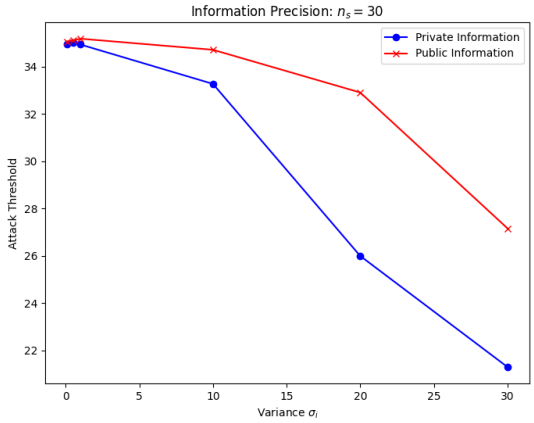
Figure 12. AI Speculator Attack Threshold and Information Accuracy (Variance) - Part I



A. Fixing Discreteness at $n_s = 10$



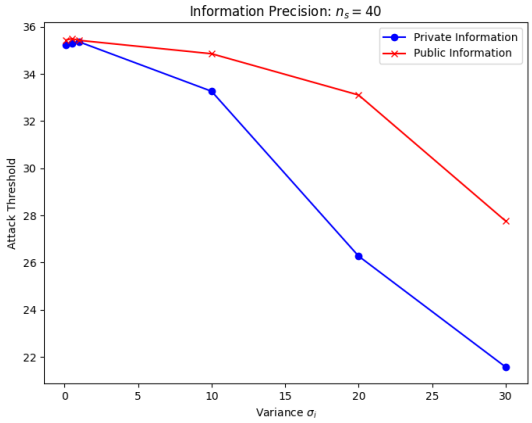
B. Fixing Discreteness at $n_s = 20$



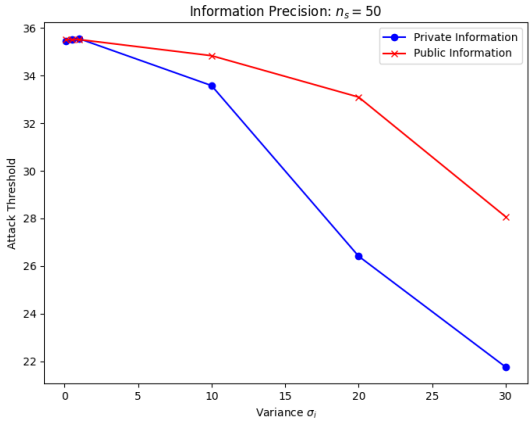
C. Fixing Discreteness at $n_s = 30$

Note: This figure compares the AI speculator attack thresholds under private and public information across different levels of information accuracy (variance, σ_i) while fixing the discreteness levels at $n_s = 10$, 20, and 30.

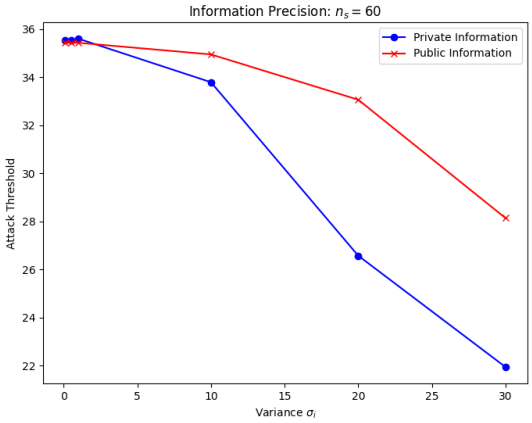
Figure 13. AI Speculator Attack Threshold and Information Accuracy (Variance) - Part II



A. Fixing Discreteness at $n_s = 40$



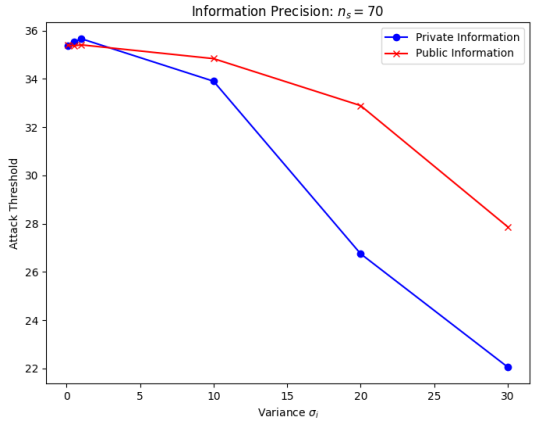
B. Fixing Discreteness at $n_s = 50$



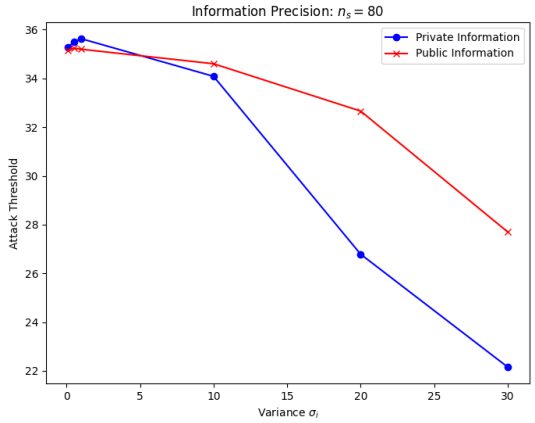
C. Fixing Discreteness at $n_s = 60$

Note: This figure compares the AI speculator attack thresholds under private and public information across different levels of information accuracy (variance, σ_i) while fixing the discreteness levels at $n_s = 40$, 50, and 60.

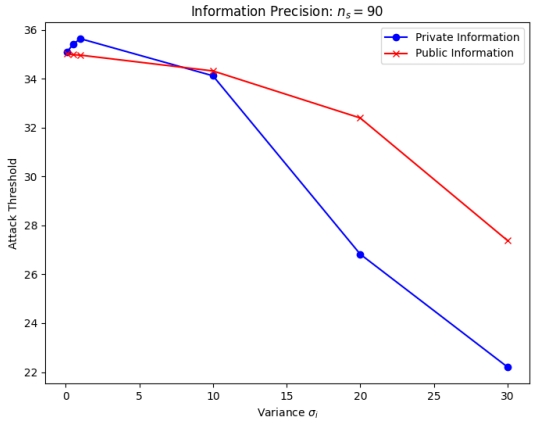
Figure 14. AI Speculator Attack Threshold and Information Accuracy (Variance) - Part III



A. Fixing Discreteness at $n_s = 70$



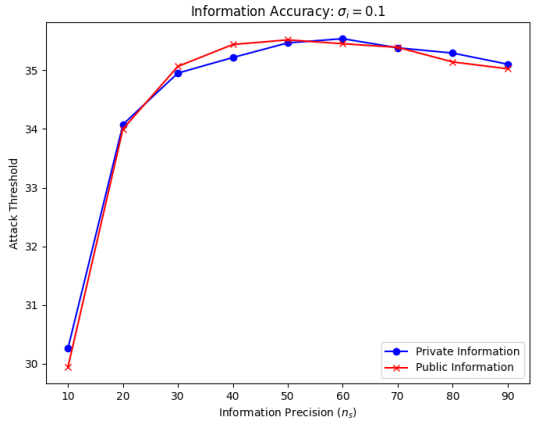
B. Fixing Discreteness at $n_s = 80$



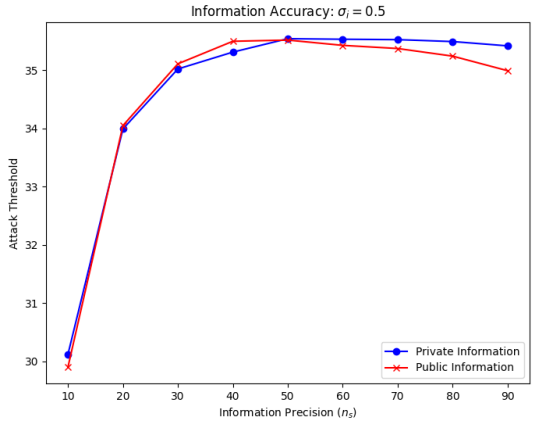
C. Fixing Discreteness at $n_s = 90$

Note: This figure compares the AI speculator attack thresholds under private and public information across different levels of information accuracy (variance, σ_i) while fixing the discreteness levels at $n_s = 70$, 80, and 90.

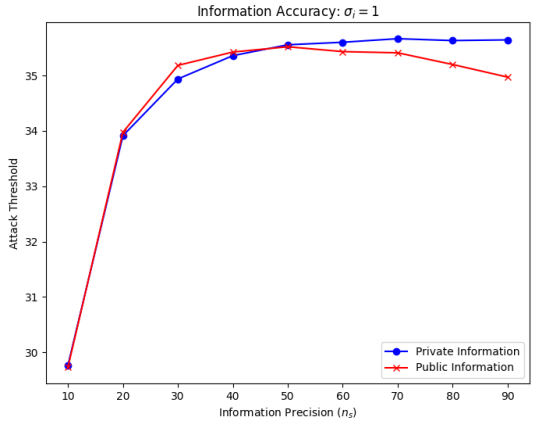
Figure 15. AI Speculator Attack Threshold and Information Precision — Part I



A. Fixing Accuracy at σ_i = 0.1



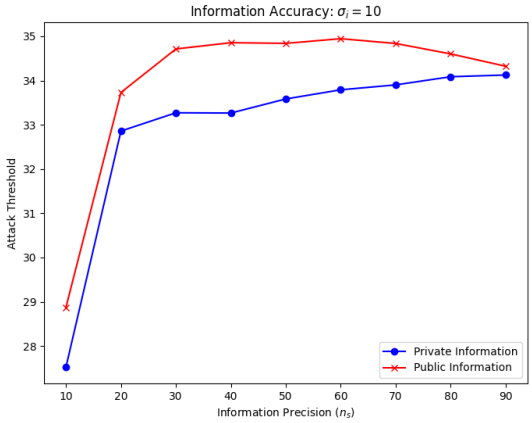
B. Fixing Accuracy at σ_i = 0.5



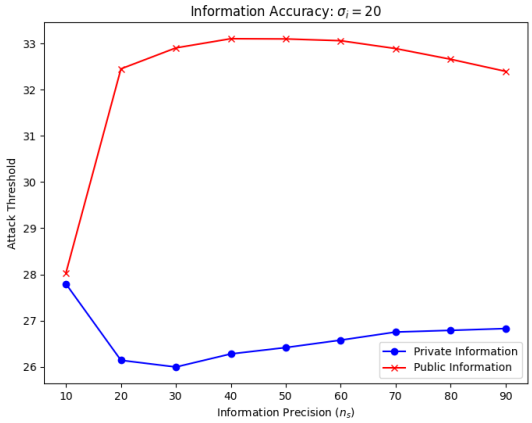
C. Fixing Accuracy at σ_i = 1

Note: This figure compares the AI speculator attack thresholds under private and public information by varying information precision (n_s) while fixing information accuracy at σ_i = 0, 1, 5, and 10 across four panels.

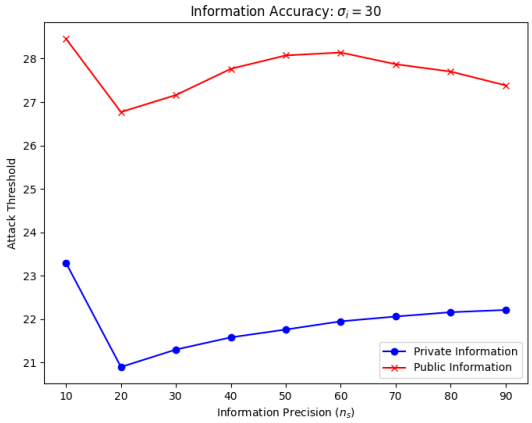
Figure 16. AI Speculator Attack Threshold and Information Precision — Part II



A. Fixing Accuracy at σ_i = 10



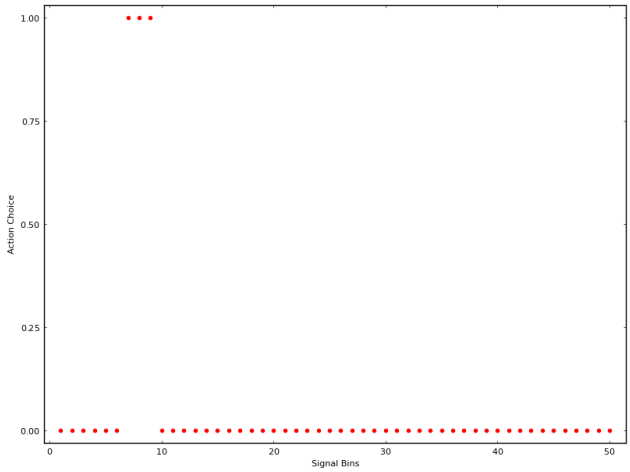
B. Fixing Accuracy at σ_i = 20



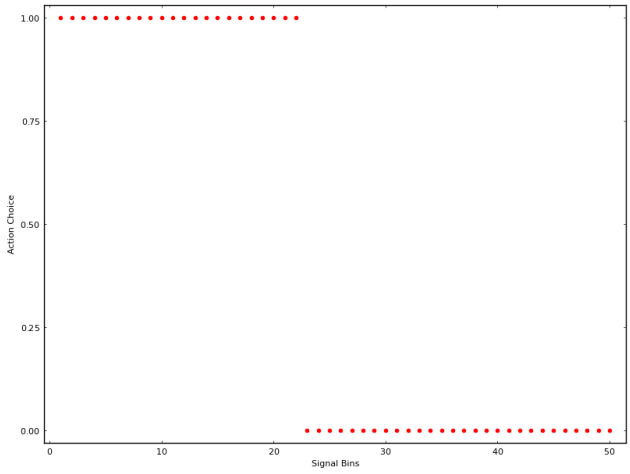
C. Fixing Accuracy at σ_i = 30

Note: This figure continues to compare AI speculator attack thresholds under private and public information by varying information precision (n_s) while fixing information accuracy at σ_i = 15, 20, 25, and 30 across four panels.

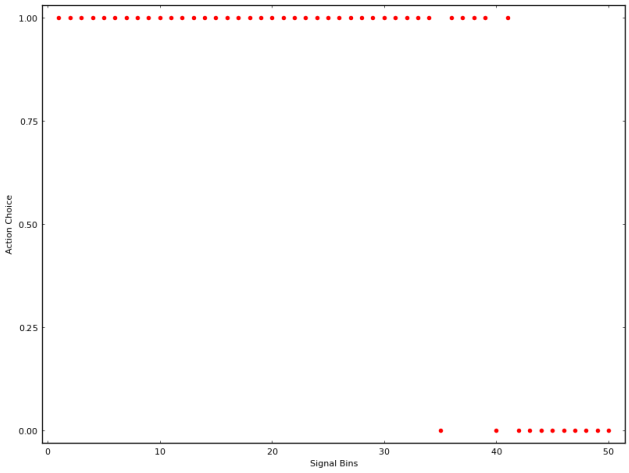
Figure 17. AI Speculators' Learned Attack Strategy with Idiosyncratic Sentiment



A. Vol. of Sentiment $\sigma_i = 51$, Discreteness $n_s = 50$



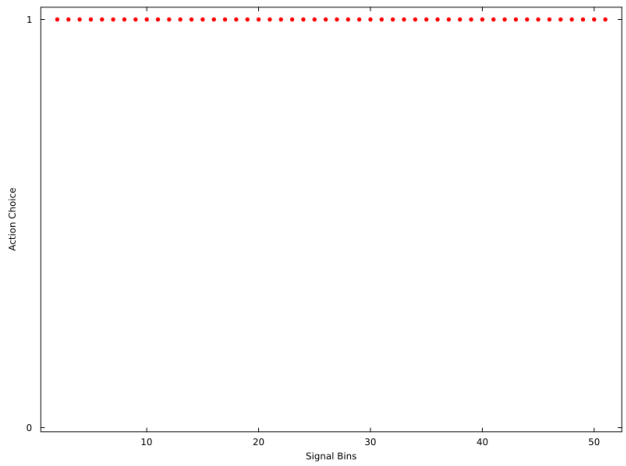
B. Vol. of Sentiment $\sigma_i = 60$, Discreteness $n_s = 50$



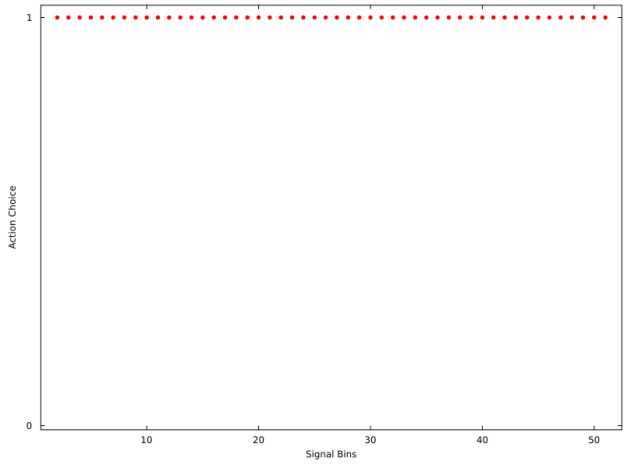
C. Vol. of Sentiment $\sigma_i = 70$, Discreteness $n_s = 50$

Note: This figure illustrates the learned attack strategy of AI speculators under varying volumes of idiosyncratic sentiment (σ_i) with fixed discreteness ($n_s = 50$). The panels represent sentiment volatility at $\sigma_i = 51$, 60, and 70, respectively. Action choice 1 means Attack, and action choice 0 means Not Attack.

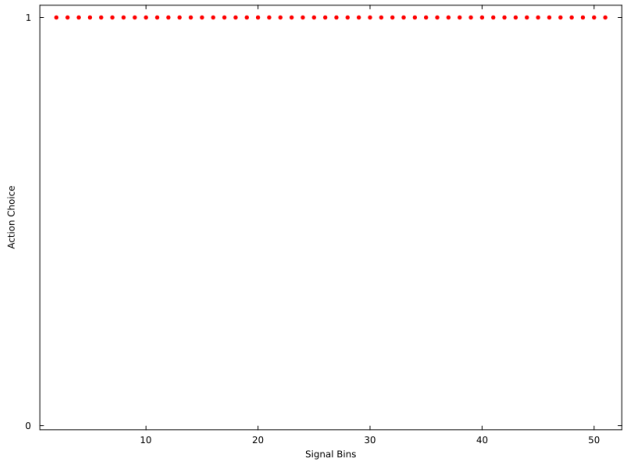
Figure 18. AI Speculators' Learned Attack Strategy with Market Sentiment



A. Vol. of Sentiment σ = 51, Discreteness n_s = 50



B. Vol. of Sentiment σ = 60, Discreteness n_s = 50



C. Vol. of Sentiment σ = 70, Discreteness n_s = 50

Note: This figure shows the learned attack strategy of AI speculators under varying volumes of market sentiment (σ) with fixed discreteness (n_s = 50). The panels depict sentiment volatility at σ = 51, 60, and 70, respectively.