AI Bubbles with Large Language Models *

Álvaro Cartea Patrick Chang Nan Chen Mingyue Zhong

This version: March 25, 2026.

ABSTRACT

AI agents participate in both rational and irrational bubbles. In a sequential bubble game, AI agents generate speculative trades even though there is a unique no-trade equilibrium. Increasing reasoning capacity reduces, but does not eliminate, such irrational bubbles. Chain-of-thought analysis shows that AI agents participate in irrational bubbles because they form incorrect beliefs through simplified expectations of the market, and they are sensitive to framing effects when forming their beliefs. When the environment admits both a no-trade equilibrium and a speculative bubble equilibrium, AI agents recognize this multiplicity and coordinate on the speculative equilibrium. Our results suggest that greater AI sophistication does not necessarily eliminate bubbles or improve market efficiency.

Keywords: Artificial Intelligence, Asset Pricing, Bubbles, LLMs, Market Efficiency.

*Álvaro Cartea is with Oxford-Man Institute, University of Oxford. Patrick Chang is with Oxford-Man Institute, University of Oxford. Nan Chen is with Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. Mingyue Zhong is with Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong. We thank the Oxford-Man Institute for funding the experiments. Nan Chen and Mingyue Zhong acknowledge support from Hong Kong Research Grant Council via the General Research Fund scheme (GRF No. 14205824 and No. 14211203)

1. Introduction

A central premise in financial economics is that competitive trading should push prices towards the fundamental value. Under the efficient market hypothesis, mispricings are difficult to exploit because the asset price reflects all available information, so prices reflect the fundamental value of the asset. Yet, speculative bubbles, episodes in which the asset price detaches and exceeds its fundamental value, remain a persistent counterexample to the efficient market hypothesis (see, e.g., Malkiel, 2003). In practice, irrational exuberance (see Shiller, 2000) in the form of speculation can lead traders to purchase overpriced assets not because fundamentals justify the price, but because they believe there are opportunities for future capital gains upon resale.

The rapid improvement and diffusion of artificial intelligence (AI) into financial decision-making raises a natural hope: if mispricings reflect limits in information processing, or limits in human cognition, then delegating decisions to AI systems that process larger datasets and that reason more consistently and at greater depth should make markets more efficient. In this paper, we study whether AI agents participate in speculative bubbles, and whether increases in AI sophistication improve market efficiency in environments prone to speculative bubbles.

We use an experimental bubble game to show that AI agents participate in both rational and irrational bubbles. In environments where equilibrium logic predicts no-trade, AI agents participate in speculative trading. Increasing the cognitive capacity of the AI agent (via higher reasoning effort) reduces, but does not eliminate, such irrational bubbles. On the other hand, when the environment admits both a no-trade equilibrium and a speculative bubble equilibrium, increasing the cognitive capacity of AI agents increases their ability to recognize the two distinct equilibria. However, conditional on recognizing both equilibria, AI agents systematically select the speculative equilibrium, regardless of their cognitive capacity. Together, these findings suggest that greater AI sophistication (in the form of cognitive capacity) mitigates inefficiencies that arise from bounded rationality, but they do not elimi-

nate inefficiencies that arise from equilibrium multiplicity; in fact, greater AI sophistication may even strengthen speculative coordination when bubbles are rational.

In this paper, we use the bubble game of Moinas and Pouget (2013, 2016) to study speculative behavior by AI agents in the form of large language models (LLMs). The bubble game is well-suited for isolating speculation because the fundamental value of the traded asset is zero, so any trade is necessarily justified by resale opportunities rather than by fundamentals. The game also cleanly separates rational and irrational bubbles depending on whether prices are capped. When the maximum feasible price is capped, the game admits a unique no-trade Bayesian Nash equilibrium, so any trade constitutes an irrational bubble.[1] When the price cap is removed, the environment can admit both a no-trade equilibrium and a speculative bubble equilibrium sustained by resale beliefs, so speculative trades can be consistent with equilibrium incentives.

We implement the bubble game as a one-shot decision problem for LLM agents. Each agent decides whether to buy an asset offered at a price of P. If the agent buys the asset, then the agent offers to resell the asset at a price of $10\,P$. If a subsequent player buys the asset, then the agent makes a profit, while failure to resell results in a loss because the fundamental value of the asset is zero, which is common knowledge. Although the underlying market is sequential, the exogenous price path (multiplying the price by 10 at each step) and the informational structure (players do not know their position in the market sequence but can infer information from the observed price) allow the game for each player to be reduced to a one-shot decision, which makes it straightforward to implement with a prompt.

Our experimental approach with AI subjects in the form of state-of-the-art reasoning models provides two methodological advantages. First, reasoning models allow us to observe their internal chain-of-thought, which, unlike elicited belief reports in human experiments, is generated as part of the inference process that directly determines the action of the agent. This allows us to study what agents do, and how they reach their decisions. Second, reasoning

[1] An infection argument in the spirit of Morris et al. (1995) implies that rational players never trade.

models provide a novel way to study the effect of bounded rationality through their reasoning effort. The reasoning effort acts as an internal thinking budget that controls the number of reasoning steps the model takes to solve the problem. Crucially, the model remains the same for different levels of reasoning effort: the underlying model weights are identical, the training data are identical, and the post-training process is identical. Adjusting the reasoning effort of the model is equivalent to adjusting the cognitive capacity of the model, where the knowledge and preferences of the model remain the same, but the difference is how the model processes and applies the knowledge it possesses.[2]

We find that AI agents systematically participate in irrational bubbles when the unique equilibrium is no-trade. The agents also exhibit the snowball effect documented in Moinas and Pouget (2013), where the propensity to buy increases as more iterated reasoning is required to derive the equilibrium strategy. Our analysis shows that increasing the cognitive capacity of the agent reduces the propensity to buy, but does not eliminate the snowball effect. Thus, the magnitude of the propensity of irrational trades is tied to the cognitive capacity of the agent, and to the cognitive demands of equilibrium reasoning.

We analyze the internal chain-of-thought of the agents to understand how their reasoning paths relate to their final decision. When agents explicitly use backward induction in their reasoning, the propensity to buy decreases, which is consistent with equilibrium logic. However, backward induction does not eliminate the snowball effect, which suggests that agents who adopt the correct approach often fail to carry the full iterated reasoning to completion. Conversely, when agents recognize that the environment is a bubble, the propensity to buy increases. That is, agents actively attempt to ride the bubble rather than exit the bubble. This pattern is difficult to reconcile with a pretraining memorization story in which agents recall the no-trade equilibrium from Moinas and Pouget (2013, 2016). Furthermore, the

[2]Experimental economics traditionally prioritizes revealed preferences over elicited reasoning or stated beliefs, in part because verbal reports may reflect ex-post rationalization or imperfect introspection. In our setting, the internal chain-of-thoughts are not elicited belief reports. Rather, they are generated as part of the inference process that determines the action of the model, and varying reasoning effort changes both the content of reasoning and the distribution of choices. We therefore use the internal chain-of-thoughts to identify mechanisms that underlie revealed behavior, not to substitute for it.

persistence of the snowball effect among agents who are aware of the bubble is consistent with speculative trades driven by beliefs about resale.

We further analyze the internal chain-of-thought of the agents to understand why they participate in irrational bubbles. Speculative motives appear in all of their reasoning, while herding-style justifications appear as a secondary reinforcement (even though herding is irrelevant to the bubble game). We map the content of their reasoning process to canonical models of behavioral game theory. We find that speculation is overwhelmingly consistent with the analogy-based expectation equilibrium of Jehiel (2005). Agents compress a complex strategic environment into a small number of analogy classes (e.g., "early" vs. "late" in the market sequence) and best respond to these simplified beliefs. This oversimplification generates incorrect beliefs about the likelihood of resale and therefore sustains speculative trades even when equilibrium logic predicts no-trade.

Next, we show that speculative behavior is highly sensitive to framing. We keep the environment fixed and change only the opponent label. When the AI agent faces AI opponents, speculation largely disappears. However, when AI agents face human opponents, speculative trading becomes almost certain for most offer prices. Furthermore, when facing human opponents, AI agents frequently adopt downward-looking beliefs consistent with cognitive hierarchy models (see, e.g., Camerer et al., 2004). These results show that framing alters the way in which AI agents form beliefs. Moreover, in mixed populations, buy rates respond nonlinearly to the population composition. A small fraction of humans can shift beliefs enough to generate a sharp increase in aggregate participation of the bubble.

Finally, we study equilibrium selection when there is no price cap and the bubble game admits both a no-trade equilibrium and a speculative bubble equilibrium. Higher reasoning effort increases the rate at which agents recognize that multiple equilibria exist. However, conditional on explicitly recognizing both equilibria, agents coordinate on the speculative equilibrium: the buy rate is 100% regardless of their reasoning level. Thus, while greater cognitive capacity mitigates inefficiencies driven by bounded reasoning in environments with

a unique no-trade equilibrium, it can increase the likelihood of speculation in environments with multiple equilibria by improving equilibria recognition and facilitating coordination on speculative outcomes.

Our results suggest that the relationship between AI sophistication and market efficiency is not monotonic. Improvements in the cognitive capacity of AI agents reduce mispricings driven by bounded reasoning, but do not necessarily eliminate bubbles. In fact, improvements in the cognitive capacity may even strengthen speculative coordination when bubbles are rational. These findings highlight the importance of understanding rational bubbles to identify environments in which improvements in AI may fail to translate into more efficient markets.[3]

Our work builds upon Moinas and Pouget (2013) and contributes to the literature that documents speculative bubbles in experimental financial markets. The pioneering study of Smith et al. (1988) documents bubbles in a double auction market for an asset that pays random dividends. Subsequent work shows that bubbles also arise when the asset has a constant fundamental value (see Noussair et al., 2001), but tend to disappear when some traders are experienced (see Dufwenberg et al., 2005), or when traders learn (see Hong et al., 2021). Other studies highlight different mechanisms behind bubble formation. Lei et al. (2001) show that speculation is not necessary for bubbles to arise, while Brunnermeier and Morgan (2010) examine how traders ride bubbles in environments with timing considerations. Our work differs from this literature by using AI agents as experimental subjects. This setting allows us to vary cognitive capacity while holding preferences and information fixed, which provides a novel way to study how bounded rationality affects bubble formation.

[3]Standard asset pricing arguments rule out bubbles under strong regularity conditions (e.g., transversality and no-arbitrage in competitive equilibrium), which implies that persistent deviations from fundamentals require departures from these assumptions (see Santos and Woodford, 1997). A large theoretical literature identifies several mechanisms through which rational bubbles can arise, including overlapping-generations economies where intrinsically unproductive assets serve as stores of value (see, e.g., Tirole, 1985; Weil, 1987); resale motives under short-sale constraints and asymmetric information (see, e.g., Allen et al., 1993; Conlon, 2004); agency and financial frictions such as delegated portfolio management and limited liability (see, e.g., Allen and Gorton, 1993; Allen and Gale, 2000); and coordination frictions among sophisticated traders (see, e.g., Abreu and Brunnermeier, 2003; Brunnermeier and Morgan, 2010; Doblas-Madrid, 2012).

More broadly, our work contributes to the growing literature on unintended consequences of AI in economic environments. A number of recent studies show that delegating decision-making to AI can generate anti-competitive outcomes. For example, algorithms can learn to collude in pricing games through reinforcement learning (see, e.g., Calvano et al., 2020; Askenazi-Golan et al., 2026), belief learning (see, e.g., Cartea et al., 2026), or even when implemented with LLMs (see, e.g., Fish et al., 2024; Keppo et al., 2026). Related work shows that LLMs can also learn predatory strategies (see, e.g., Nikandrova and Parekh, 2025). Closest to our setting is Dou et al. (2025) who show that reinforcement learning algorithms can learn collusive trading strategies that create and exploit (negative) bubbles in financial markets. Our contribution differs from this literature by studying speculative behavior in a controlled experimental market and by examining how variation in cognitive capacity affects bubble formation and equilibrium selection among AI agents.

The remainder of the paper proceeds as follows. Section 2 describes the bubble game and the experimental design with AI agents. Section 3 presents the main results on bubble participation. Section 4 analyzes the reasoning of the agents and the mechanisms that underlie speculative behavior. Section 5 studies framing effects and mixed populations. Section 6 examines equilibrium recognition and selection when speculative equilibria exist. Section 7 concludes.

2. Experimental Setup

We use the bubble game of Moinas and Pouget (2013, 2016) to study whether AI agents will participate in speculative bubbles. Below, we outline and analyze the bubble game, and we outline the experimental protocol of our AI subjects.

2.1. Bubble Game

The bubble game is an experimental market designed to study speculative bubbles. Even though the bubble game features a sequential market for an asset, the design of the game allows one to reduce the game to a one-shot game, which makes it extremely simple to implement with AI agents in the form of LLMs.

2.1.1. Game Structure

The bubble game is a market game where players sequentially trade an asset that does not generate any income and has a fundamental value of zero that is common knowledge. The game consists of N players that is common knowledge, and each player is randomly assigned a position in the trading sequence with equal probability. Players do not know their position in the trading sequence. However, players can infer some information upon observing the price at which they are offered because the price path is exogenous and evolves by powers of 10.

The game proceeds sequentially. The first player is offered the opportunity to buy the asset at a price of $P_1 = 10^{j-1}$, where j follows a (truncated) geometric distribution with parameter 1/2. That is, $\mathbb{P}(j = k) = (1/2)^k / c$ for $k \in \{1, \ldots, K\}$, where $c = \sum_{k=1}^{K} (1/2)^k$ is a normalizing constant. If the first player buys the asset, then he proposes to resell the asset to the second player at a price 10 times higher, i.e., $P_2 = 10\, P_1$. Continuing this sequence, if the n^{th} player decides to buy the asset at P_n, then he proposes to resell the asset to the next player at a price $P_{n+1} = 10\, P_n = 10^n\, P_1$. At any point of the game, if a player is unable to resell the asset, then the game ends.

Each player is endowed with one dollar. Their decision is either to buy the asset or not to buy the asset. If the asset price is greater than one dollar, players can still buy the whole asset, but upon reselling, players only receive the proportion of the payoffs according to how much they invested, i.e., $1/P_n$.[4]

[4] Moinas and Pouget (2013, 2016) include an outside financier in their setup. However, the outside

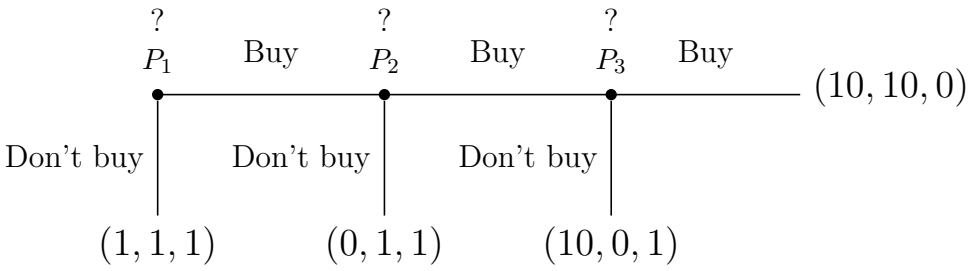


Fig. 1: Timing of the bubble game with $N = 3$. The bubble game features a sequential market. The question marks indicate that players are equally likely to be in any position of the trading sequence, and to be offered to buy at prices P_1, P_2, or P_3.

Consider a player who buys the asset at a price of P_n. If he is unable to resell the asset, then his final wealth is zero because the fundamental value of the asset is zero. On the other hand, if he is able to resell the asset, then he receives $1/P_n$ percent of the resale value of $P_{n+1} = 10\,P_n$. Thus, his final wealth is 10 dollars. Finally, if a player refuses to buy the asset, or if the game ends before his turn, then he retains his initial endowment of one dollar. Figure 1 illustrates the timing of the bubble game with $N = 3$ and the associated payoffs.

2.1.2. Equilibrium Analysis

The bubble game features both rational and irrational bubbles depending on whether or not there is a price cap. For each player i, speculating is rational if

$$\begin{aligned} &[1 - \mathbb{P}\,(\text{last})] \cdot \mathbb{P}\,(\text{next player buys}) \cdot u_i\,(10) \\ &+ \big[\,[1 - \mathbb{P}\,(\text{last})] \cdot \mathbb{P}\,(\text{next player does not buy}) + \mathbb{P}\,(\text{last})\,\big] \cdot u_i\,(0) \geq u_i\,(1) \end{aligned} \tag{IRC}$$

holds, where $u_i\,(\,\cdot\,)$ is the utility of player i, and $\mathbb{P}\,(\text{last})$ is the posterior probability that player i is the last player conditional on the price he is offered.

The equilibrium depends on whether there is a cap on the maximum price. With finite players N, if we truncate the geometric distribution so that $P_1 = 10^{K-1}$ is the highest possible realization of the first price, then this implies that the maximum price of the game is capped at 10^{K+N-2}. On the other hand, if we allow K to go to infinity, so the distribution

financier is not included in their experiments, so we only present the setup without the outside financier.

of the first price P_1 is geometric, then there is no (implicit) cap on the maximum price.

Price Cap When K is finite, there is a unique Bayesian Nash equilibrium in which a bubble never forms because all the players do not trade as a consequence of a backward induction type argument.

In the bubble game, the maximum price cap may not be reached, so players can never fully know they are the last in the market sequence unless they are offered the maximum price cap. As a consequence, standard backward induction cannot start. Instead, the equilibrium is solved through an infection argument in the spirit of Morris et al. (1995). Consider a player who is offered the asset at the maximum price cap. The player understands that he is last in the market sequence with no possibility to resell the asset. Thus, he refuses to buy.

Next, consider a player who is offered one price below the maximum price cap. Upon observing the price, the player infers that he is either last or second to last in the market sequence. If he is last in the market sequence, then he refuses to buy. Similarly, if he is second to last, he anticipates the next player will also refuse upon offered the maximum price cap. Thus, he also refuses to buy. This iterated reasoning process continues all the way to all of the lower prices, so the unique equilibrium is no bubble. In general, a player who is offered a price of P needs to iterate this reasoning for $\log_{10}\left(10^{K+N-2}/P\right)$ steps to derive the equilibrium strategy. As a consequence, if players buy the asset when there is a price cap, then an irrational bubble forms.

No Price Cap When the value of K goes to infinity so that the distribution of the first price is geometric, a rational bubble can arise as an equilibrium if (IRC) holds for all players. The bubble equilibrium exists because players are never sure to be the last in the market sequence. Indeed, if players anticipate that other players will speculate, then there exist increasing and concave utility functions $u_i(\cdot)$ for which (IRC) holds.

To see this in a general setting, consider the set of potential prices defined as $P \in \{m^n \text{ for } m > 1 \text{ and } n \in \mathbb{N}\}$, where m is the price explosiveness parameter, which we set as

$m = 10$, then the posterior probability of being the last player conditional on the price is

$$\mathbb{P}\left(\text{last}\right) = \begin{cases} 0 & \text{if } P = m^n < m^{N-1}\,, \\ \frac{1-q}{1-q^N} & \text{if } P = m^n \geq m^{N-1}\,, \end{cases} \tag{1}$$

where P is the offered price, and $q \in (0,1)$ is the parameter of the geometric distribution of the first price, which we set as $q = 1/2$.

If we use (1) to rewrite (IRC), it is straightforward to show that for risk-neutral players, there exists a bubble equilibrium if and only if

$$m \geq \frac{1-q^N}{q-q^N}\,.$$

This condition is less restrictive as the number of players N increases, or if the price is more explosive, i.e., the value of m is larger. Similarly, suppose players have constant relative risk aversion utility given by $u_i\left(x\right) = x^{1-\gamma_i}/(1-\gamma_i)$ if $\gamma_i > 0, \gamma_i \neq 1$, and given by $u_i\left(x\right) = \ln(x)$ if $\gamma_i = 1$, where γ_i is the level of risk aversion for player i. For $\gamma_i \neq 1$, there exists a bubble equilibrium if and only if

$$\gamma_i \leq 1 - \frac{\ln\left(\frac{1-q^N}{q-q^N}\right)}{\ln\left(m\right)}$$

for all $i \in \{1, \ldots, N\}$. This condition must hold for all players, and becomes less restrictive if there are more players, or if the price is more explosive.[5] Thus, if there is no price cap and if (IRC) holds for all players, then a rational bubble forms when players buy the asset because the asset trades at a value that exceeds its fundamental value.

Finally, no bubble is always an equilibrium. If players anticipate that other players will reject the offer, then they are better off if they refuse to buy the asset.

[5] In both the risk-neutral and the constant relative risk aversion utility, we focus on when the posterior probability of being last player is non-zero because (IRC) trivially holds otherwise. For a more comprehensive analysis, see the Supplementary Material of Moinas and Pouget (2013) where they also discuss mixed strategies.

2.2. AI Subjects

We implement the bubble game with LLMs as our experimental subjects. The LLMs take a textual instruction ("prompt") as an input and provide a textual response as the output. In the context of the bubble game, this is similar to the protocol with human subjects, where written instructions are given to human subjects and their responses are recorded. In our experiments, we use the state-of-the-art reasoning model of GPT-5.2 under two configurations: low reasoning effort and high reasoning effort. The benefit of reasoning models is that one can extract and analyze their reasoning process.

2.2.1. LLMs and Reasoning Models

Fundamentally, LLMs are neural networks that are trained to predict the next token in a sequence of text, where a token is a sequence of short characters. Most modern LLMs use the transformer architecture of Vaswani et al. (2017) that allows the neural network to incorporate the context surrounding a token in a cheap and scalable fashion.

State-of-the-art models such as OpenAI's GPT-5 are trained in two distinct phases. First, the model is pretrained through supervised learning from diverse datasets that include public and licensed text (see OpenAI, 2025). The objective in the pretraining phase is to predict the next token. This phase of training is where the model acquires its knowledge base.

The second phase of post-training differs for reasoning and non-reasoning models. The post-training objective for non-reasoning models is to maximize expected human satisfaction, and this is trained through reinforcement learning with human feedback (see OpenAI, 2023). On the other hand, the post-training objective for reasoning models is to train the model to reason through reinforcement learning (see OpenAI, 2025).

The success of reasoning models stems from chain-of-thought prompting. Wei et al. (2022) show that generating a series of intermediate reasoning steps improves the ability of LLMs to perform complex reasoning. Through reinforcement learning, reasoning models learn to tackle a problem through a series of intermediate reasoning steps. This allows the

model to refine their thinking process, to try different strategies, and to recognize their mistakes (see OpenAI, 2025). Consequently, reasoning models are trained to think before they answer. They generate an internal chain-of-thought which one can extract and analyze to understand their approach to solving the problem.[6] This phase of training is where the model learns to use its knowledge base.

Reasoning models can be configured to control the number of reasoning tokens they use during inference. This acts as an internal thinking budget that controls the number of reasoning steps the model takes to solve the problem. For example, with low effort, the model may prematurely end its reasoning steps and return to the first reasonable solution; while with high effort, it may extend its reasoning steps, backtrack if it finds an error, and verify the answer before responding.

Crucially, the model remains the same for different levels of reasoning effort: the underlying model weights are identical, the training data are identical, and the post-training process is identical. Adjusting the reasoning effort of the model is equivalent to adjusting the cognitive capacity of the model, where the knowledge of the model remains the same, but the difference is how the model processes and applies the knowledge it possesses. With reasoning models, one can study the effect of bounded rationality, in the form of cognitive capacity, on bubble formation in AI subjects.

2.2.2. Prompt Structure

By design, the bubble game can be reduced to a one-shot game. Thus, the bubble game is played by simply prompting our subjects. There are three components to our prompts.

1. **System Prompt.** The agent is tasked to play a game. Here, we follow the best practices for reasoning models and encourage the model to reason freely.
2. **Game Information.** The details of the game are outlined with examples and the payoffs are spelt out. We modify the instructions of Moinas and Pouget (2013, 2016)

[6] OpenAI does not expose the raw reasoning tokens from the model for safety reasons and for competitive secrecy. Instead, they provide a summary of the model's reasoning.

so that the parameters of the game can be easily adjusted.

3. **Treatment.** The agent is given the parameters of the game along with any treatment specific details.

Appendix A provides the details for all the prompts we use with concrete examples.

In all prompts, we restrict the subject to respond with either a YES or a NO. The internal chain-of-thought from the model is extracted directly from the API.[7]

2.3. Experimental Design

An important consideration for our experiments is that LLMs are not deterministic. Specifically, they sample from a probability distribution to generate the next token. As a consequence, this can lead to different responses in the same experimental setting.[8] This stochasticity also extends to their internal chain-of-thought, which can result in different approaches to solving the same problem.

Formally, chain-of-thought (CoT) models can be represented sequentially as

$$\begin{aligned}
\mathrm{CoT}_1 &\sim p\left(\mathrm{CoT}_1 \,|\, \text{prompt}, \text{reasoning}\right) \\
\mathrm{CoT}_2 &\sim p\left(\mathrm{CoT}_2 \,|\, \text{prompt}, \text{reasoning}, \mathrm{CoT}_1\right) \\
&\vdots \\
\mathrm{CoT}_n &\sim p\left(\mathrm{CoT}_n \,|\, \text{prompt}, \text{reasoning}, \mathrm{CoT}_{1:n-1}\right) \\
a &\sim p\left(a \,|\, \text{prompt}, \text{reasoning}, \mathrm{CoT}_{1:n}\right) ,
\end{aligned}$$

where $a \in \{\mathrm{YES}, \mathrm{NO}\}$ is the final decision of the model. The length of the realization of the chain-of-thought path $\mathrm{CoT}_{1:n}$ is random but depends on the reasoning level. Thus, the

[7] The default setting of reasoning models do not include the output from the reasoning summary, so this parameter must be explicitly activated.

[8] In previous versions, such as GPT-4, one can control the level of stochasticity through the temperature parameter. However, this feature is not available for GPT-5.2.

choice rule of the LLM is the conditional distribution

$$p\left(a \mid \text{prompt}, \text{reasoning}\right) = \int p\left(a \mid \text{prompt}, \text{reasoning}, \text{CoT}\right)\, p\left(\text{CoT} \mid \text{prompt}, \text{reasoning}\right)\, d\text{CoT}\,,$$

where the reasoning level alters the distribution over possible chain-of-thought paths and therefore the induced distribution over final answers.[9] To address this stochasticity, we repeat each treatment 100 times to characterize both their behavior and their reasoning process.

In the baseline experiments, the number of players is $N = 100$, and we consider offer prices $P \in \{10, 100, 10^{10}, 10^{50}, 10^{90}\}$, price cap parameter $K \in \{5, 10, 20, 100, 200\}$, and reasoning levels $r \in \{\text{low}, \text{high}\}$. This produces a $5 \times 5 \times 2$ factorial design, and each treatment is given by the triplet (P, K, r). Here, the unique equilibrium is no-trade because the price is capped, so trades in this setting is participation in irrational bubbles.

3. Bubble Participation

We study bubble participation of AI agents and document three core findings. First, reasoning models systematically participate in speculative bubbles and exhibit the snowball effect: the buy rate increases as more iterated reasoning is required to derive the equilibrium strategy, i.e., the buy rate increases as $\log_{10}\left(10^{K+N-2}/P\right)$ increases. Second, increasing the reasoning effort reduces bubble participation. Third, AI agents are less likely to participate in the bubble if they use backward induction in their reasoning process, but they are also more likely to participate in the bubble if they recognize the game as a bubble.

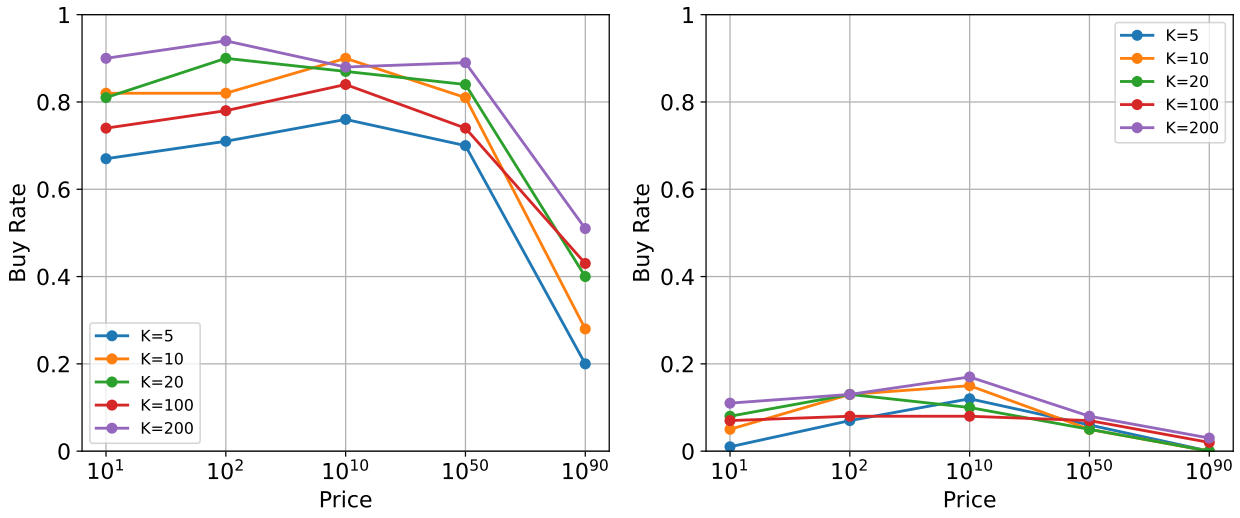
Economically, this establishes two implications for market outcomes. First, AI agents participate in irrational bubbles. That is, AI agents can generate substantial speculative trades in environments where a backward induction argument implies a unique no-trade equilibrium. Second, the magnitude of the propensity of speculative trades is tied to the

[9] In contrast, the choice rule for non-reasoning models without internal chain-of-thought is $p\left(a \mid \text{prompt}\right)$.

cognitive capacity of the AI agent, and the cognitive demands of equilibrium reasoning. This ties the formation of irrational bubbles to the bounded rationality of AI agents.

3.1. Cognitive Capacity and the Snowball Effect

We examine the buy rates of the AI agents. For each treatment triplet (P, K, r), we estimate the buy rate as the fraction of YES decisions out of the 100 repetitions. Figure 2 shows the buy rates as a function of the offered prices for each price cap parameter K, and for low and high reasoning.



(a) Low reasoning (b) High reasoning

Fig. 2: **Buy rate for low and high reasoning**.

Several patterns emerge. First, bubble participation is economically significant. With low reasoning, buy rates are often between 70-90% across a wide range of offered prices and price cap parameters. Second, increasing the cognitive capacity of the AI agent reduces their participation in speculative trades. However, even with high reasoning, non-trivial buy rates persist (around 10-20%) for a wide range of offered prices and price cap parameters.

This shows that state-of-the-art AI agents participate in speculative trades even when equilibrium logic predicts no-trade because the price is capped. Economically, this means

that AI agents are willing to hold an overpriced asset which can only be justified by resale opportunities, not fundamentals; that is, AI agents participate in speculation.

Finally, AI agents also exhibit the snowball effect documented in Moinas and Pouget (2013, 2016). Specifically, the buy rate increases as more iterated reasoning is required to derive the equilibrium strategy, i.e., the buy rate increases as $\log_{10}\left(10^{K+N-2}/P\right)$ increases.[10] This corresponds to a higher buy rate for lower offer prices P and for higher price caps K. The implication is that AI agents are more likely to generate bubbles in environments that require deeper iterated reasoning.

Table 1: **Wilcoxon signed-rank test p-values**.

(a) Pairwise comparison of the offer price P.

	$P = 100$	$P = 10^{20}$	$P = 10^{50}$	$P = 10^{90}$
$P = 10$	0.0076	0.0098	1.0000	0.0020
$P = 100$	·	0.2357	0.0020	0.0020
$P = 10^{20}$	·	·	0.0059	0.0020
$P = 10^{50}$	·	·	·	0.0020
$P = 10^{90}$	·	·	·	·

(b) Pairwise comparison of the price cap K.

	$K = 10$	$K = 20$	$K = 100$	$K = 200$
$K = 5$	0.0109	0.0208	0.0137	0.0020
$K = 10$	·	0.4990	0.2754	0.0127
$K = 20$	·	·	0.1055	0.0075
$K = 100$	·	·	·	0.0020
$K = 200$	·	·	·	·

Observations are matched at the treatment level. Each pairwise test is based on 10 matched treatments.

To formalize these patterns, we use the Wilcoxon signed-rank test to determine whether the median of the differences between paired observations is significantly different from zero. Table 1 reports the p-values from the Wilcoxon signed-rank test for pairwise comparisons of the offer price P (top panel) and pairwise comparisons of the price cap K (bottom panel).

[10] Another element of the snowball effect is that the buy rate increases as the probability of not being last increases. For the combinations of offered prices and price caps we consider, the probability of being last is always zero, so we do not investigate this element of the snowball effect.

The observations are matched at the treatment level. Thus, for each pair of offer prices, we match buy rates within each fixed (K, r) combination and treat them as paired observations; similarly, for each pair of price caps, we match buy rates within each fixed (P, r) combination and treat them as paired observations. This leads to 10 matched samples for each pairwise comparison.

The results support the snowball effect. Pairwise comparisons against $P = 10^{90}$ are all significant with p-values of 0.002. Cross-referencing with Figure 2 shows that the buy rates for lower offer prices are statistically greater than the buy rate for $P = 10^{90}$. Similarly, pairwise comparisons against $K = 200$ are all significant at the 1% level. Thus, the buy rate for $K = 200$ is statistically greater than the buy rates for lower price caps. Therefore, lower offer prices and higher price caps lead to increased buy rates, which is direct support for the snowball effect.

Finally, we quantify the effect of cognitive capacity by comparing buy rates under low and high reasoning with observations matched at the treatment level. The difference is significant with a p-value of $5.96 \cdot 10^{-8}$, which shows that increasing the cognitive capacity of the AI agents reduces bubble participation. This, together with the snowball effect, connects the bounded rationality of AI agents to the formation of irrational bubbles.

We estimate a binomial generalized linear model with a logit link at the treatment level to sharpen our statistical analysis. Let p_t denote the buy probability in the treatment triplet $t = (P, K, r)$. We estimate

$$\text{logit}\,(p_t) = \beta_0 + \beta_1\,\text{Steps}_t + \beta_2\,\mathbb{1}_{\text{Reasoning = High}} + \beta_3\,\text{Steps}_t \times \mathbb{1}_{\text{Reasoning = High}} + \alpha_K\,,$$

where $\text{Steps}_t = K + N - 2 - \log_{10}(P)$ is the number of reasoning steps required to derive the equilibrium strategy under the treatment triplet $t = (P, K, r)$, and $\mathbb{1}_{\text{Reasoning = High}} = 1$ when the reasoning level is high and is zero when the reasoning level is low. Finally, α_K are fixed effects for the price cap.

Table 2: **Snowball Effect Regression**.

	Coefficient	p-value
Intercept	-1.1226^{***}	0.000
Steps	0.0226^{***}	0.000
$\mathbb{1}_{\text{Reasoning = High}}$	-3.6587^{***}	0.000
Steps $\times$ $\mathbb{1}_{\text{Reasoning = High}}$	-0.0017	0.141
Observations	50	

Binomial generalized linear model with a logit link estimated at the treatment level. Fixed effects for K included but not reported. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

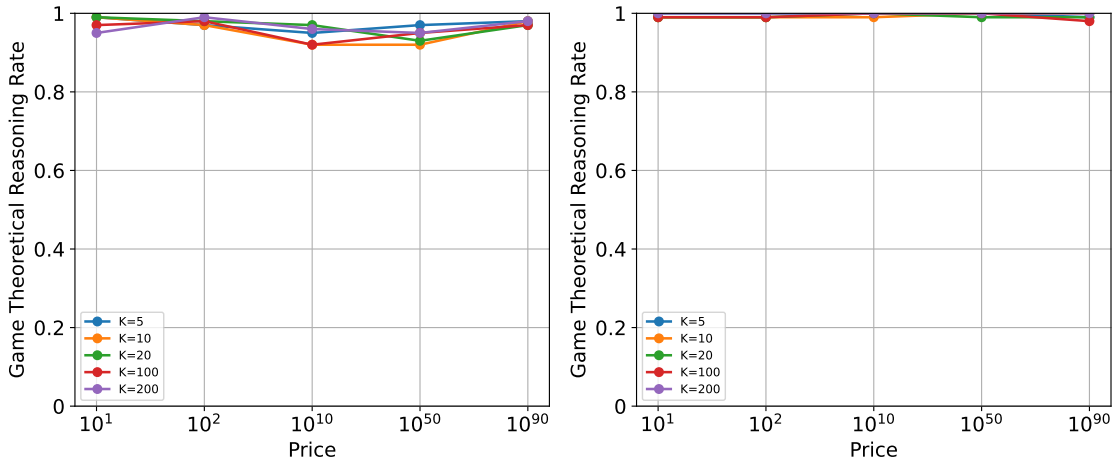
Table 2 reports the results of the regression and confirms the patterns in Figure 2. The positive and significant coefficient on Steps confirms the snowball effect: the propensity to buy increases with the cognitive demands of equilibrium reasoning. Furthermore, increasing the reasoning level, i.e., the cognitive capacity of the AI agent, reduces the propensity to buy. However, the interaction term is negative but not statistically significant. Therefore, higher reasoning reduces the propensity to buy, but does not eliminate the snowball effect.

While the decisions of the AI agents alone are insightful, their internal chain-of-thought provides valuable insights into their underlying logic and latent reasoning structures. In particular, we can study their approach to solving a problem (i.e., their procedural rationality, see Simon, 1957) and how it relates to their final decision.

3.2. Procedural Rationality and Chain-of-Thought

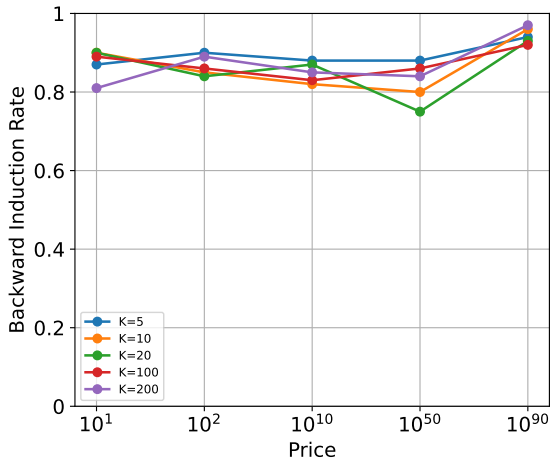
Our results show that deviations from the no-trade equilibrium depends on the difficulty of equilibrium reasoning and the cognitive capacity of the AI agents. Here, we use the internal chain-of-thought to measure the procedural rationality of the AI agents and relate it to bubble participation.

For each realization, we classify whether the agent (i) uses game theoretical reasoning to solve the problem, (ii) uses backward induction to solve the problem, and (iii) is aware that the game is a bubble. Figure 3 shows the rates of these reasoning features as a function of

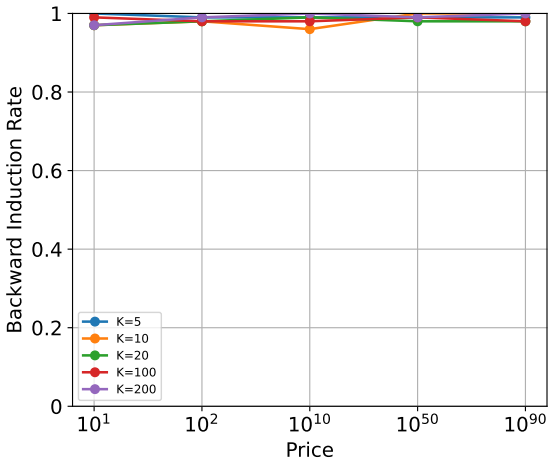


(a) Low reasoning

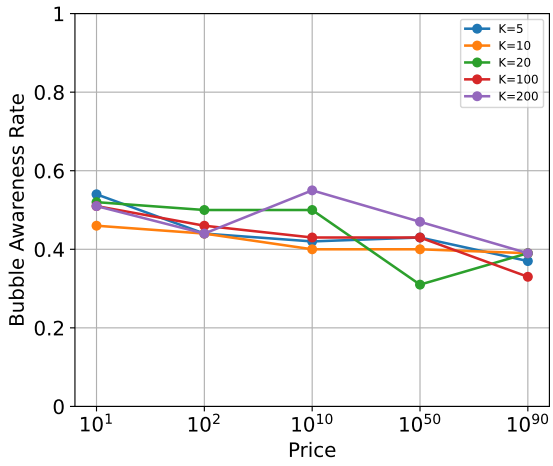
(b) High reasoning



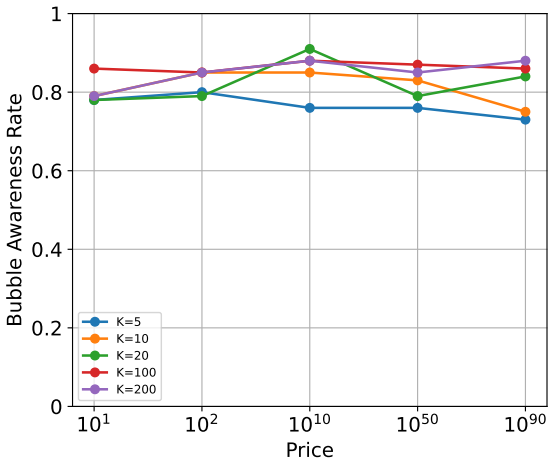
(c) Low reasoning



(d) High reasoning

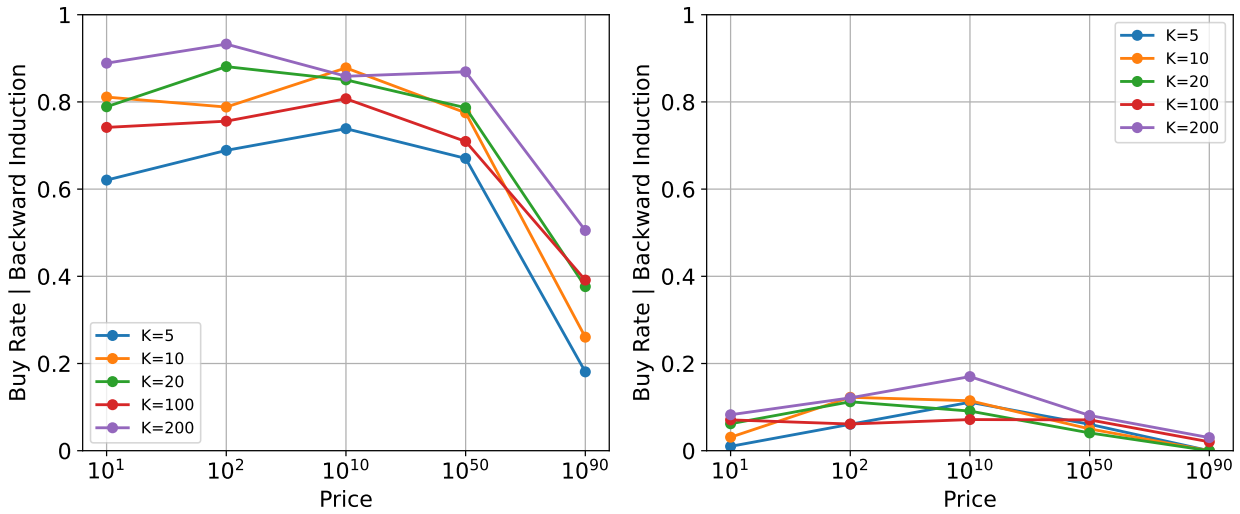


(e) Low reasoning



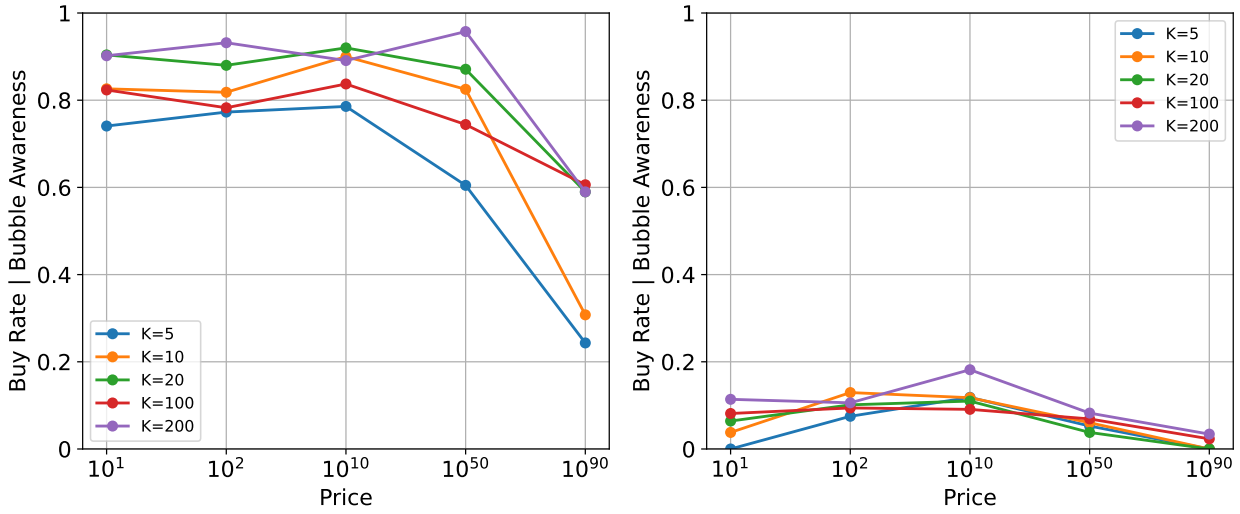
(f) High reasoning

Fig. 3: **Rates of Game Theoretical Reasoning, Backward Induction, and Bubble Awareness**.



(a) Low reasoning (b) High reasoning

Fig. 4: **Buy Rate Conditional on Backward Induction**.



(a) Low reasoning (b) High reasoning

Fig. 5: **Buy Rate Conditional on Bubble Awareness**.

the offered prices for each price cap parameter K, and for low and high reasoning. High reasoning increases the prevalence of each category, particularly backward induction and bubble awareness. The Wilcoxon signed-rank test shows that the increase in each category due to high reasoning are all significant at the 1% level. This shows that increasing the reasoning level increases the procedural rationality of the agent: it increases the probability that they use equilibrium-relevant reasoning when solving the problem.

Figures 4 and 5 condition the buy rates on backward induction and bubble awareness, respectively. The conditioning separates two economically distinct channels.

First, conditioning on backward induction highlights the equilibrium reasoning channel. Theoretically, if the agent applies backward induction correctly, then the bubble unravels and no-trade is the unique outcome. However, Figure 4 shows that backward induction does not eliminate irrational bubbles. Instead, the persistence of the snowball effect highlights the cognitive limitations of the agents. Thus, the agents may understand the structure of the problem, but fail to carry through the required iterated reasoning to solve the problem.

Second, conditioning on bubble awareness highlights the speculation channel. Even though recognizing the environment as a bubble does not pin down the no-trade outcome, Figure 5 shows that the agents actively ride the bubble.

To test these channels formally, we construct binomial counts for realizations with and without a given reasoning feature within each treatment. We use binomial generalized linear models with a logit link at the treatment and CoT-feature level to estimate

$$\text{logit}\left(p_{t,f}\right) = \beta_0 + \beta_1 \, \mathbb{1}_{f \in \text{CoT}} + \beta_2 \, \text{Steps}_t + \beta_3 \, \mathbb{1}_{\text{Reasoning = High}} + \alpha_K \, ,$$

where $\mathbb{1}_{f \in \text{CoT}}$ denotes the channel we study, and $f \in \{\text{Backward Induction}, \text{Bubble Awareness}\}$ is the reasoning feature. Thus, $\mathbb{1}_{\text{Backward Induction} \in \text{CoT}} = 1$ if the agent uses backward induction in his reasoning process and is zero otherwise. Similarly, $\mathbb{1}_{\text{Bubble Awareness} \in \text{CoT}} = 1$ if the agent recognizes the game as a bubble in his reasoning process and is zero otherwise.

For each channel, we also estimate the model with and without the interaction between Steps $\times \mathbb{1}_{f \in \text{CoT}}$, which results in a total of four models.

Table 3: **Backward Induction and Bubble Awareness on the Propensity to Buy**.

	Backward Induction		Bubble Awareness	
	(1)	(2)	(3)	(4)
Intercept	1.1298***	0.9869*	−1.1776***	−1.2046***
	(0.006)	(0.089)	(0.000)	(0.001)
$\mathbb{1}_{\text{Backward Induction} \in \text{CoT}}$	−2.3503***	−2.2018***		
	(0.000)	(0.000)		
$\mathbb{1}_{\text{Bubble Awareness} \in \text{CoT}}$			0.2582**	0.3136
			(0.017)	(0.167)
Steps	0.0221***	0.0233***	0.0221***	0.0223***
	(0.000)	(0.000)	(0.000)	(0.000)
$\mathbb{1}_{\text{Reasoning = High}}$	−3.8275***	−3.8277***	−4.0209***	−4.0156***
	(0.000)	(0.000)	(0.000)	(0.000)
Steps $\times \mathbb{1}_{\text{Backward Induction} \in \text{CoT}}$		−0.0013		
		(0.757)		
Steps $\times \mathbb{1}_{\text{Bubble Awareness} \in \text{CoT}}$				−0.0004
				(0.746)
Observations	95	95	100	100

Binomial generalized linear models with a logit link. Fixed effects for K included but not reported. Standard errors clustered at the treatment level. p-values in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 reports the results of the regressions with standard errors clustered at the treatment level.[11] The table shows two key findings.

First, when the agent uses backward induction, the propensity to buy decreases. This effect is statistically significant and highlights the effect of procedural rationality. The result suggests that bubbles partly persist because the agents do not always take the correct approach to solving the problem. However, backward induction does not eliminate the snowball effect: the coefficient of the interaction term is close to zero and statistically insignificant. Together, the implication is that while backward induction reduces the propensity to buy,

[11] For high reasoning, there are five treatments where the rate of backward induction is 100%. Our results do not change if we restrict the samples to treatments where within treatment comparisons exist.

agents who use backward induction also fail to carry through with the required iterated reasoning to solve the problem.

Second, recognizing the game as a bubble does not eliminate speculation. Instead, the propensity to buy increases. This contrast also addresses a concern about pretraining. Our experiment is based on the bubble game of Moinas and Pouget (2013, 2016), so the concern is that most likely the agent has been exposed to their paper during pretraining.[12] Therefore, the concern is that the agent recalls the no-trade equilibrium from its training data. If their decision was driven by memorizing the equilibrium solution, then recognizing the game as a bubble should decrease participation. Instead, we find the opposite. This pattern is difficult to reconcile with a memorization argument because the agents do not regurgitate the no-trade solution. Instead, they continue to speculate upon recognizing the bubble.

Finally, the persistence of the snowball effect among agents who are aware of the bubble is consistent with speculative trades driven by beliefs about resale. The next section investigates the formation of beliefs that underlie their speculative behavior, and connects their beliefs and decisions to models of behavioral game theory.

4. AI Reasoning

We examine the mechanism that underlies speculation by AI agents, and why the agents deviate from the no-trade equilibrium when the price of the asset is capped. An equilibrium requires players to have beliefs that are consistent with the actual strategies played, and that each player chooses a best response based on these beliefs. Therefore, do AI agents speculate because they form correct beliefs, but do not best respond (see, e.g., Quantal Response Equilibria of McKelvey and Palfrey, 1995, 1998)? Or because AI agents best respond, but form incorrect beliefs (see, e.g., Camerer et al., 2004; Jehiel, 2005)?

A key benefit with reasoning models is that we can analyze the reasoning summaries to study the justifications that accompany speculative trades. Instead of fitting structural

[12] GPT-5.2's knowledge cutoff date is 31 August, 2025.

models to the observed behavior to infer which behavioral model best fits the data (see, e.g., Moinas and Pouget, 2013), we directly observe their reasoning process through their internal chain-of-thought. This allows us to move beyond behavioral inference and examine which behavioral model best matches the reasoning process.

Our analysis shows that AI agents participate in speculation because they oversimplify the environment when forming their beliefs. Specifically, the agents bundle the nodes of the game into analogy classes consistent with the analogy-based expectation equilibrium (ABEE) of Jehiel (2005).

4.1. Reasoning Themes

We study the high-level themes that the AI agents use to justify their decision to participate in the irrational bubble. We restrict our attention to realizations in which the agent participates in the bubble, and he also recognizes the environment as a bubble.

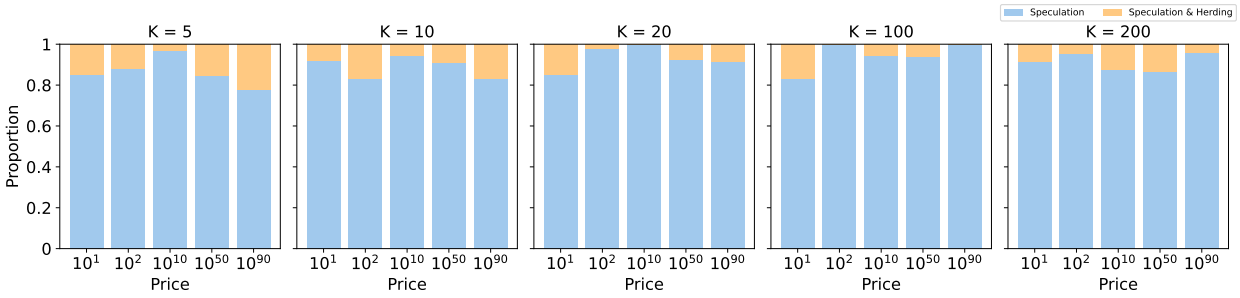
Within this subset, we classify the reasoning summaries into two themes:

1. Speculation: the agent tries to ride and to profit from the bubble because he believes the next person is likely to buy.
2. Herding: the agent uses the actions of previous players to inform his decision to buy.

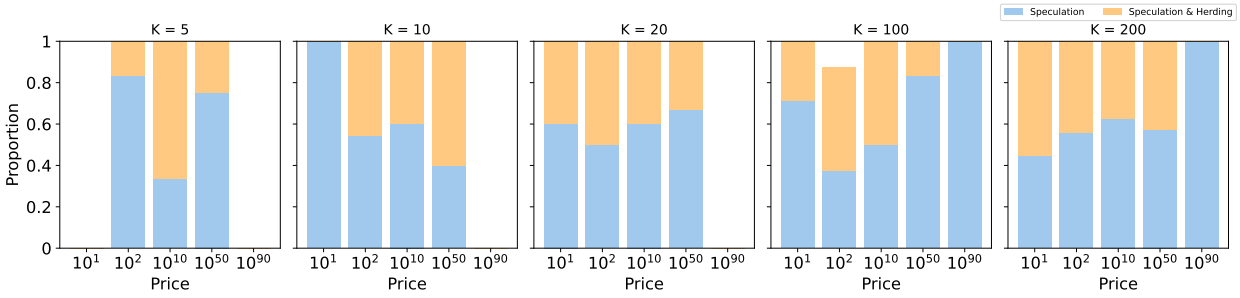
We allow multiple themes within a single realization. An agent can have both speculative motives and simultaneously justify their decision through the actions of previous players. Although the bubble game is not designed to study herding, we consider herding as a reasoning theme because, despite playing no role in the equilibrium logic of the game, it frequently emerges in the internal chain-of-thought of the agent.[13]

Figure 6 shows the classification based on the two reasoning themes. The classification shows that speculation is the dominant reasoning theme. Speculative motives appear in all

[13]The bubble game is not designed to study herding because it does not have the usual forces that generate rational herding. For example, it does not have information externalities present in the herding model of Banerjee (1992), or the information cascade model of Bikhchandani et al. (1992). Nor does it have the reputational concerns of Scharfstein and Stein (1990), or coordination problems with multiple equilibria of Diamond and Dybvig (1983) when there is a price cap.



(a) Low Reasoning



(b) High Reasoning

Fig. 6: **Reasoning Themes**.

Classification is performed on answers where the agent participates in the bubble, and also recognizes the environment as a bubble. Prices with no bars are because no samples exist for the treatment triplet.

the internal chain-of-thoughts.[14] Although herding arguments appear with a non-trivial frequency, the herding arguments always appear along with a justification based on speculative motives. This suggests that bubble participation is primarily driven by beliefs about the potential for future resale, while herding acts as a secondary justification to reinforce the decision of the agent to participate in the irrational bubble.

4.2. Belief Formation under Speculation

To gain further insights into why AI agents speculate, we analyze which behavioral model best matches the reasoning process of the agent. Following Moinas and Pouget (2013), three behavioral models are most relevant:

[14]The bar in the bottom panel of Figure 6 does not sum to 100% because the internal chain-of-thought was incomplete, so the agent did not justify his answer.

1. The ABEE of Jehiel (2005). The agent oversimplifies the problem into analogy classes, and plays a best response based on the belief about the behavior in each analogy class. Two analogy classes are relevant here: (i) players who are sure of not being last, so they are early in the bubble, and (ii) players who have a non-zero probability of being last, so they are late in the bubble.

We classify the reasoning process as ABEE if the agent buys because he inferred that he, or the next player, is early or not too late in the market sequence.

2. The cognitive hierarchy (CH) model of Camerer et al. (2004). The agent holds downward-looking beliefs about his opponents.

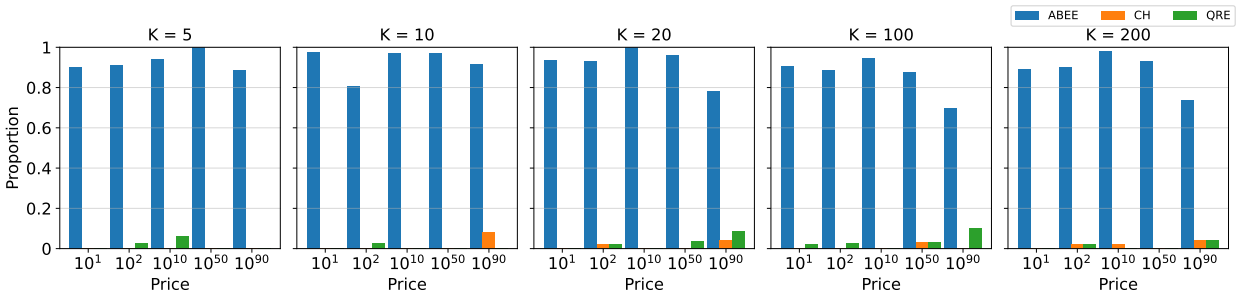
We classify the reasoning process as CH if the agent buys because he believes that his opponents are less sophisticated than himself.

3. Quantal response equilibria (QRE) of McKelvey and Palfrey (1995, 1998). The solution concept postulates that players hold correct beliefs but make errors when best responding. Players do not always choose the best action, but choose what is better more often.

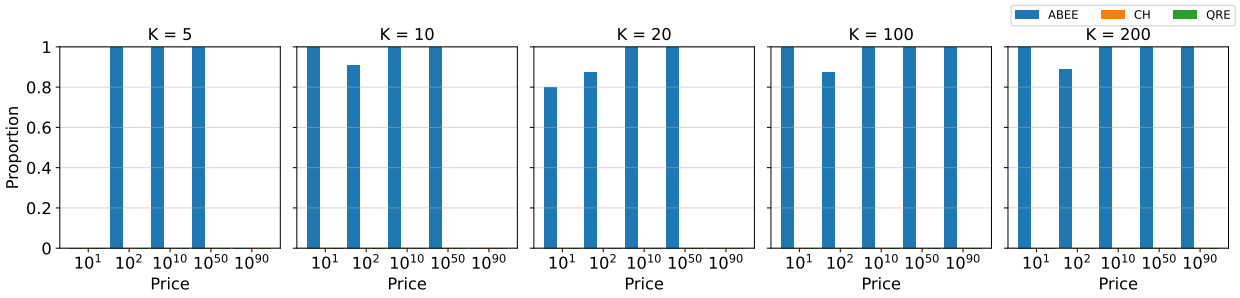
We classify the reasoning process as QRE if the final decision of the agent is inconsistent with his reasoning. Our classification for QRE is weaker than the definition because we do not require the agent to have correct beliefs.[15]

Figure 7 shows the classification based on the three behavioral models. The classification shows that ABEE is the most common reasoning process that leads to speculation. This reasoning process holds across both low and high reasoning, and the ABEE reasoning accounts for over 80% of the answers across treatment triplets. This also explains why recognizing

[15] At a mechanical level, the process LLMs use to stochastically generate the next token resembles a quantal response: tokens are generated probabilistically rather than deterministically. However, they key aspect of QRE is the payoff responsiveness of the final action. Thus, keeping with the spirit of QRE, our classification focuses on whether the reasoning is consistent with the final decision.



(a) Low Reasoning



(b) High Reasoning

Fig. 7: **Speculative Beliefs.**

Classification is performed on answers where the agent participates in the bubble, and also recognizes the environment as a bubble. Prices with no bars are because no samples exist for the treatment triplet.

the bubble does not reduce participation. Recognizing the situation as a bubble does not deter participation if the agent finds it profitable to speculate given his incorrect belief.

The broader implication is that AI agents can often oversimplify difficult strategic problems. Rather than working through the full chain of reasoning required to determine the equilibrium outcome, AI agents form simplified expectations of the market. Consequently, AI agents sustain irrational bubbles because they compress a complex strategic environment into a tractable mental shortcut. This compression of strategic complexity can contribute to the persistence of bubbles in markets where decisions are delegated to AI agents, and if these agents also rely on similar simplifications. For example, if some investors follow this ABEE approach, then Bianchi and Jehiel (2025) show that bubbles and crashes can arise in a purely speculative market in which information is complete and it is commonly understood that the bubble cannot grow forever.

5. Opponent Framing and Belief Formation

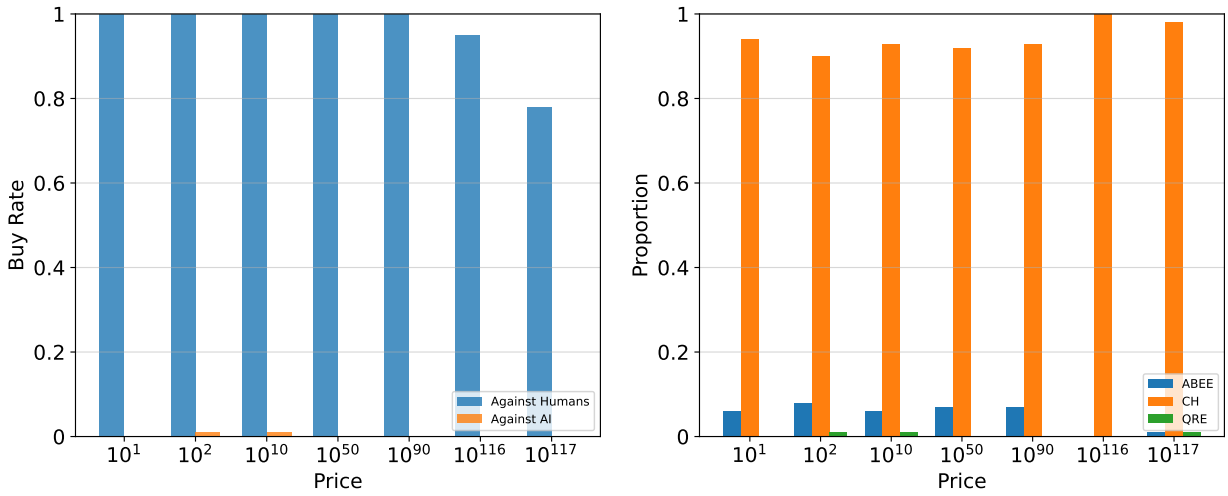
We examine whether the behavior of the AI agent depends on the perceived type of opponent. This question is economically relevant for markets that adopt the use of AI agents. If AI agents condition their beliefs on the perceived sophistication of their counterparties, then market outcomes may depend not only on incentives but also on the composition of the population.

We find that opponent labels materially affect behavior. When facing human opponents, the AI agents form downward-looking beliefs consistent with CH models. When facing AI opponents, speculation largely disappears. Our results show that belief formation in AI agents depends on the opponents they face, and that a small change in market composition can significantly alter speculative dynamics.

5.1. Opponent Type

We study the effect of framing based on the type of opponent the AI agent faces. We focus only on high reasoning, and we fix $K = 20$ and vary the offer price $P \in \{10, 100, 10^{10}, 10^{50}, 10^{90}, 10^{116}, 10^{117}\}$ to include offer prices close to the price cap of 10^{118}. For each offer price, we run the prompt with two opponent labels: (i) "opponents are other independent copies of yourself", and (ii) "opponents are humans". All other parts of the prompt remain identical.

Figure 8a plots the buy rate by the offer price for each type of opponent. The figure shows that the buy rates are highly sensitive to opponent labels. When facing AI opponents, speculation largely disappears. On the other hand, when facing human opponents, speculative trading is almost certain for all offer prices other than the penultimate offer price. The results show that AI agents condition their behavior on the perceived sophistication of their opponents. That is, the same strategic environment can produce vastly different outcomes depending on the label of the opponent. This resembles the effect of framing in Tversky and



(a) Buy Rate

(b) Decomposition of Speculation against Humans

Fig. 8: **Buy Rate and Speculative Beliefs by Opponent Type**.
Classification is performed on answers where the agent participates in the bubble.

Kahneman (1986). The implication here is that market stability can be influenced by informational framing effects in AI systems. Specifically, framing effects can amplify or dampen speculative dynamics, independent of underlying incentives.

A closer look at the buy rate against humans at the penultimate offer price shows that, in many instances, AI agents perceive human opponents as level-0 players, who act randomly and non-strategically. In fact, for speculation to be profitable here, the human opponent must be unable to infer that being offered the asset at price $P = 10^{118}$ means that he is last in the market sequence with probability one. The results show that AI agents can exploit the bounded rationality of humans, but at the same time, AI agents can also underestimate the rationality of humans, potentially to their own detriment (we like to believe that most humans are more sophisticated than level-0 players).

To confirm this downward-looking belief against humans, Figure 8b classifies the reasoning process against humans based on the three behavioral models. The classification shows that CH is the most common reasoning process that leads to speculation against humans opponents; that is, AI agents speculate because they believe humans are less sophisticated

than themselves. The results show that the framing effect alters the way in which AI agents forms beliefs.

5.2. Population Mix

We analyze how small changes in the framing effect can drastically alter market outcomes. We fix $K = 20$, $P = 10^{10}$, and consider population compositions where the fraction of humans lies between $[0, 0.15]$, and the remaining population consists of AI agents.

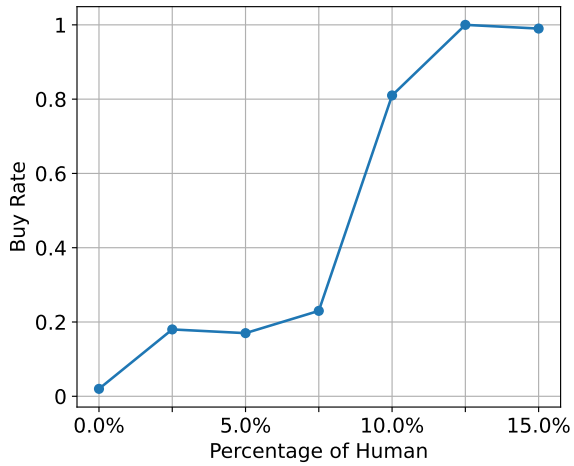


Fig. 9: **Buy Rate for Various Proportions of Human Opponents**.

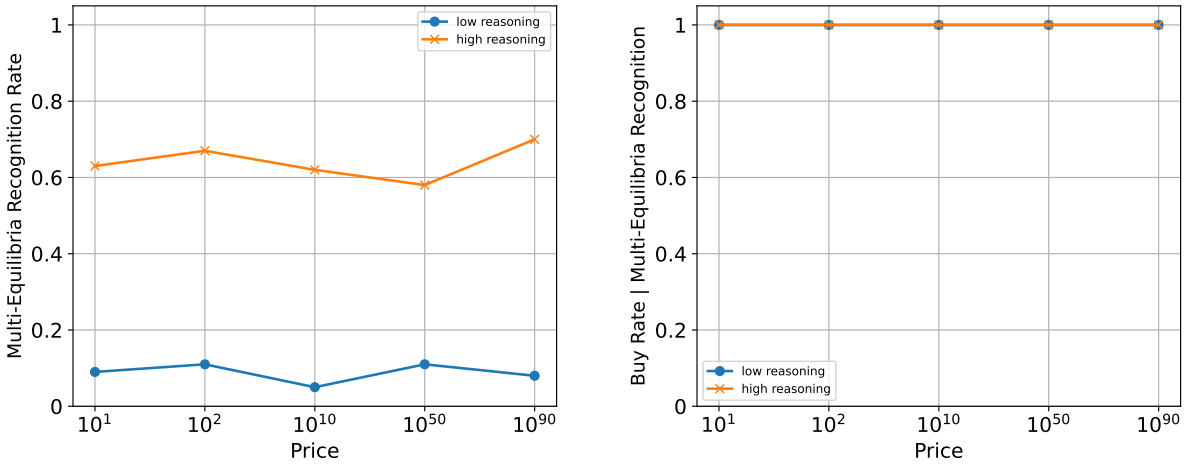
Figure 9 shows the buy rate as a function of the percentage of humans in the population composition. The figure shows a non-linear sigmoid-like relationship where the buy rate is near zero when the population is all AI agents, and then increases sharply as the human population grows. The results show that a small fraction of human players can alter the belief of AI agents in a way that materially changes aggregate participation.

Although manipulating the label of the opponent is a framing effect, our results have implications for market efficiency. In the bubble game with a price cap, the no-trade equilibrium is the efficient market benchmark, and deviations from the efficient price arises through speculation. When agents face other AI agents, speculative trade is substantially lower and outcomes are closer to efficient markets. In mixed environments, where even a small frac-

tion of participants are human, speculation increases and deviations from efficient markets become larger. Thus, market efficiency in AI-driven markets may depend on population composition. Purely AI markets and mixed human-AI markets can generate systematically different outcomes, not because fundamentals differ, but because belief formation differs.

6. Equilibrium Selection

We return to the baseline experiments and analyze how AI agents select between equilibria. When there is no price cap, i.e., $K = \infty$, the bubble game admits a no-trade equilibrium and a rational bubble equilibrium. We study whether AI agents recognize the two distinct equilibria, and if they do, which equilibrium do they select.



(a) Multi-Equilibria Recognition Rate

(b) Buy Rate Conditional on Multi-Equilibria Recognition

Fig. 10: **Multi-Equilibria Recognition and Conditional Buy Rates with no Price Cap.**

Figure 10a analyzes the internal chain-of-thought and visualizes how frequently the agents recognize that the game admits multiple equilibria when there is no price cap. The results show that equilibria recognition rates are substantially higher with high reasoning, and a Wilcoxon signed-rank test shows that this increase in equilibria recognition is statistically

significant at the 10% level (p-value = 0.0625). This shows that the ability of the agent to recognize the two equilibria depends on their cognitive capacity.

Figure 10b visualizes the buy rate conditional on recognizing the two equilibria. The results show that whenever agents explicitly acknowledge both the no-trade and the bubble equilibrium, the conditional buy rate is 100% across reasoning levels. Together, the results show that cognitive capacity affects equilibrium recognition, and, conditional on recognizing both equilibria, AI agents systematically select the speculative equilibrium, regardless of cognitive capacity.

The implication is that, when given a choice between the no-trade equilibrium and the speculative bubble equilibrium sustained by resale beliefs, AI agents systematically select the speculative equilibrium instead of the risk dominant no-trade equilibrium. The broader economic implication is that while increasing the cognitive capacity of AI agents reduces irrational bubbles in environments with a unique equilibrium, it also increases the likelihood of speculation in rational bubbles by raising the probability that agents recognize and coordinate on speculative equilibria when multiple equilibria exist.

7. Discussion

The central finding of this paper is that AI agents participate in both rational and irrational bubbles.

In environments where bubbles are irrational, AI agents generate substantial speculative activity. Increasing the cognitive capacity of the AI agents reduces speculation, but does not eliminate irrational bubbles. The internal chain-of-thought of the agents reveals that speculation in irrational bubbles is systematic. Rather than executing the full chain of iterated reasoning required to solve the equilibrium, AI agents participate in irrational bubbles because they form incorrect beliefs through simplified expectations of the market, which is consistent with the analogy-based expectation equilibrium of Jehiel (2005).

Furthermore, AI agents are sensitive to framing effects when forming their beliefs. Even when the strategic environment is fixed with a unique no-trade equilibrium, changing only the opponent label leads to substantial differences in speculative behavior. When facing human opponents, AI agents form downward-looking beliefs consistent with cognitive hierarchy models and speculate because they expect humans to act less strategically. However, when facing other AI agents in the same environment, speculative trades sharply decline. Therefore, small changes in framing alters how AI agents form beliefs, which materially affects market outcomes.

On the other hand, when the strategic environment admits both a no-trade equilibrium and a speculative bubble equilibrium sustained by resale beliefs, increasing the cognitive capacity of the AI agents increases their ability to recognize the two distinct equilibria. However, conditional on recognizing both equilibria, AI agents systematically select the speculative equilibrium, regardless of cognitive capacity. This suggests that when speculation is supported by equilibrium incentives, AI agents are willing to coordinate on speculative outcomes.

These results highlight an important distinction. When bubbles are irrational, speculation reflects bounded reasoning; but when bubbles are rational, speculation reflects equilibrium coordination. Increasing the cognitive capacity of AI agents changes the type of bubble participation, but it does not necessarily eliminate bubbles. More capable agents produce fewer irrational bubbles, but remain willing to engage in rational bubbles when the environment permits it.

These findings have direct implications for market efficiency. The relationship between the capability of AI agents and market efficiency is not monotonic. Improvements in AI agents mitigate inefficiencies that arise from bounded rationality, but they do not eliminate inefficiencies that arise from equilibrium multiplicity.

References

Abreu, Dilip and Markus K. Brunnermeier (2003): "Bubbles and Crashes," *Econometrica*, 71, 173–204.

Allen, Franklin and Douglas Gale (2000): "Bubbles and crises," *Economic Journal*, 110, 236–255.

Allen, Franklin and Gary Gorton (1993): "Churning Bubbles," *Review of Economic Studies*, 60, 813–836.

Allen, Franklin, Stephen Morris, and Andrew Postlewaite (1993): "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, 61, 206–229.

Askenazi-Golan, Galit, Domenico Mergoni Cecchelli, Edward Plumb, and Clemens Possnig (2026): "The Bounds of Algorithmic Collusion; Q-learning, Gradient Learning, and the Folk Theorem," *Available at arXiv 2411.12725*.

Banerjee, Abhijit V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Bianchi, Milo and Philippe Jehiel (2025): "Bubbles and Crashes with Partially Sophisticated Investors," *Mathematics and Financial Economics*, 19, 735–755.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Brunnermeier, Markus K. and John Morgan (2010): "Clock Games: Theory and Experiments," *Games and Economic Behavior*, 68, 532–550.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicoló, and Sergio Pastorello (2020): "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review*, 110, 3267–97.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119, 861–898.

Cartea, Álvaro, Patrick Chang, José Penalva, and Harrison Waldon (2026): "Algo-

rithmic Collusion and a Folk Theorem from Learning with Bounded Rationality," *Games and Economic Behavior*, 157, 1–21.

Conlon, John R. (2004): "Simple Finite Horizon Bubbles Robust to Higher Order Knowledge," *Econometrica*, 72, 927–936.

Diamond, Douglas W. and Philip H. Dybvig (1983): "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 91, 401–419.

Doblas-Madrid, Antonio (2012): "A Robust Model of Bubbles with Multidimensional Uncertainty," *Econometrica*, 80, 1845–1893.

Dou, Winston Wei, Itay Goldstein, and Yan Ji (2025): "Financial Market Fragility in the Era of AI Planning," *Available at SSRN 5763222*.

Dufwenberg, Martin, Tobias Lindqvist, and Evan Moore (2005): "Bubbles and Experience: An Experiment," *American Economic Review*, 95, 1731–1737.

Fish, Sara, Yannai A. Gonczarowski, and Ran I. Shorrer (2024): "Algorithmic Collusion by Large Language Models," *Available at arXiv 2404.00806*.

Hong, Jieying, Sophie Moinas, and Sébastien Pouget (2021): "Learning in Speculative Bubbles: Theory and Experiment," *Journal of Economic Behavior & Organization*, 185, 1–26.

Jehiel, Philippe (2005): "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory*, 123, 81–104.

Keppo, Jussi, Yuze Li, Gerry Tsoukalas, and Nuo Yuan (2026): "On the Fragility of AI Agent Collusion," *Available at SSRN 5386338*.

Lei, Vivian, Charles N. Noussair, and Charles R. Plott (2001): "Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica*, 69, 831–859.

Malkiel, Burton G. (2003): "The Efficient Market Hypothesis and Its Critics," *Journal of Economic Perspectives*, 17, 59–82.

McKelvey, Richard D. and Thomas R. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 6–38.

——— (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1, 9–41.

Moinas, Sophie and Sébastien Pouget (2013): "The Bubble Game: An Experimental Study of Speculation," *Econometrica*, 81, 1507–1539.

——— (2016): "The Bubble Game: A Classroom Experiment," *Southern Economic Journal*, 82, 1402–1412.

Morris, Stephen, Rafael Rob, and Hyun Song Shin (1995): "p-Dominance and Belief Potential," *Econometrica*, 145–157.

Nikandrova, Arina and Anushree Parekh (2025): "Algorithmic Exclusion by Large Language Models," *Available at SSRN 5550761*.

Noussair, Charles, Stephane Robin, and Bernard Ruffieux (2001): "Price Bubbles in Laboratory Asset Markets with Constant Fundamental Values," *Experimental Economics*, 4, 87–105.

OpenAI (2023): "GPT-4 Technical Report," *Available at arXiv 2303.08774*.

——— (2025): "OpenAI GPT-5 System Card," *Available at arXiv 2601.03267*.

Santos, Manuel S. and Michael Woodford (1997): "Rational Asset Pricing Bubbles," *Econometrica*, 65, 19–57.

Scharfstein, David S. and Jeremy C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80, 465–479.

Shiller, Robert J. (2000): *Irrational Exuberance*, Princeton, NJ: Princeton University Press.

Simon, Herbert A. (1957): *Models of Man*, New York: Wiley.

Smith, Vernon L., Gerry L. Suchanek, and Arlington W. Williams (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119–1151.

Tirole, Jean (1985): "Asset Bubbles and Overlapping Generations," *Econometrica*, 53, 1499–1528.

Tversky, Amos and Daniel Kahneman (1986): "Rational Choice and the Framing of Decisions," *Journal of Business*, 59, S251–S278.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin (2017): "Attention is All you Need," in *Advances in Neural Information Processing Systems*, ed. by I. Guyon, U. Von Luxburg, S. Bengio,

H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Curran Associates, Inc., vol. 30.

Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou (2022): "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA: Curran Associates Inc., NIPS '22.

Weil, Philippe (1987): "Confidence and the Real Value of Money in an Overlapping Generations Economy," *Quarterly Journal of Economics*, 102, 1–22.

Appendix A. Prompts

System Prompt: Your task is to play a game and you should carefully analyze the problem before answering. You may reason freely and thoroughly.

Game Information: This game consists of N players. Each player is given 1 dollar that can be used to buy the asset. In this game, your decision is to choose whether or not to buy the asset. This asset does not generate any income. If the asset price is greater than 1 dollar, you can still buy the whole asset, but when reselling, you receive the proportion of the profits according to how much you invested and your initial capital.

The game proceeds sequentially. The first player is offered the opportunity to buy at a price P_1. If he buys the asset, he proposes to sell the asset to the second player at a price 10 times higher ($P_2 = 10\,P_1$). If the second player buys the asset, then the first player collects a total of 10 dollars. The second player then proposes to sell the asset to the third player at a price 10 times higher ($P_3 = 10\,P_2 = 10^2\,P_1$). If the third player buys the asset, then the second player collects a total of 10 dollars. Continuing this logic, if the n^{th} player buys the asset, then the $n^{th} - 1$ player collects a total of 10 dollars. The n^{th} player then proposes to sell the asset to the $n^{th} + 1$ player at a price 10 times higher ($P_{n+1} = 10\,P_n = 10^n\,P_1$). Finally, if the N^{th} player buys the asset, then the $N^{th} - 1$ player collects a total of 10 dollars. The N^{th} player cannot find anyone to sell the asset to, so if he buys the asset, then he loses

1 dollar because the asset has a fundamental value of 0 dollars. Similarly, if the n^{th} player decides not to buy, then the $n^{th} - 1$ player loses 1 dollar. If the n^{th} player decides not to buy, then the game ends and all remaining players keep their dollar.

You do not know your position in the game. Positions are randomly determined with equal probability. The price proposed to the first player, P_1 is random and given by the following probabilities: $P_1 = 10^{j-1}$ with probability $p_j = (1/2)^j/c$ for $j \in \{1, \cdots, K\}$, and $c = \sum_{j=1}^{K}(1/2)^j$. Know that the proposed prices can give you information about your position in the game if you use Bayesian updating.

Treatment (example): Now let $K = 5$, i.e., P_1 has K possible values. There are $N = 100$ players, and you are offered a price of 10^5. Do you buy the asset or not? Answer only YES or NO.