

SPECULATIVE OVERPRICING IN ASSET MARKETS WITH INFORMATION FLOWS

BY THOMAS R. PALFREY AND STEPHANIE W. WANG¹

In this paper, we derive and experimentally test a theoretical model of speculation in multiperiod asset markets with public information flows. The speculation arises from the traders' heterogeneous posteriors as they make different inferences from sequences of public information. This leads to overpricing in the sense that price exceeds the most optimistic belief about the real value of the asset. We find evidence of speculative overpricing in both incomplete and complete markets, where the information flow is a gradually revealed sequence of imperfect public signals about the state of the world. We also find evidence of asymmetric price reaction to good news and bad news, another feature of equilibrium price dynamics under our model. Markets with a relaxed short-sale constraint exhibit less overpricing.

KEYWORDS: Asset pricing, heterogeneous beliefs, speculation, experimental finances.

1. INTRODUCTION

THIS PAPER STUDIES equilibrium pricing dynamics in a simple dynamic asset market where traders have heterogeneous beliefs and face the short-sale constraint. We analyze a model that follows from a long line of theoretical research initiated by Harrison and Kreps (1978; henceforth HK). That line of research has had a major impact in the theoretical finance literature, so it is surprising that there have been no attempts to directly observe one of the central implications of the theory, what we refer to as *speculative overpricing*. By speculative overpricing, we refer to the phenomenon where the current price of an asset exceeds the maximum amount any trader is willing to pay if he/she has to hold the asset to maturity (overpricing). Traders are willing to “overpay” in equilibrium because they believe (correctly) that in equilibrium there is a chance that another trader will value the asset more highly than they do at some future date. The key insight of the seminal HK paper is that speculative overpricing of a multiperiod asset can arise in equilibrium if there is a combination of the short-sale constraint and divergent beliefs about the fundamentals determining the underlying value of the asset. We report the results of a laboratory study that implements the main features of such asset markets. The transactions data from these markets are then used to test the speculative overpricing hypothesis as well as several other testable implications of the model.

¹We gratefully acknowledge the financial support of the National Science Foundation (Grant SES-0617820), the Gordon and Betty Moore Foundation, the Social Science Experimental Laboratory at Caltech, and the California Social Science Experimental Laboratory at UCLA. We are grateful for comments from Peter Bossaerts, Ron Harstad, Tony Kwasnica, Stephen Morris, Howard Rosenthal, three referees, a co-editor, and participants at numerous seminars and conferences.

The model is by design a simple one, simple enough to study easily in the laboratory using the standard multiple-unit open-book continuous double-auction market. We assume a finite horizon, two equally likely states of the world, A and B , and a single asset, a simple Arrow–Debreu security that yields a payoff of 1 in state A and of 0 in state B . As with most of the literature following HK, traders are assumed to be risk-neutral. In each time period, a new public information signal arrives at the market that is observed by all traders. Signals are binary and independent and identically distributed (i.i.d.), conditional on the state. The source of belief heterogeneity is motivated by well-documented heterogeneity in how individuals update prior beliefs after receiving a signal that is correlated with the state of the world.² Specifically, some individuals over-react to signals in the sense of updating their prior beliefs more sharply than would a Bayesian, while other individuals under-react in the sense of updating their prior beliefs more conservatively than would a Bayesian. If the traders are drawn from a pool of over-reacters and under-reacters, then the posterior beliefs of traders can differ even after observing the same sequence of public signals.

Together, the short-sale constraint and heterogeneous beliefs result in higher equilibrium prices than in the case where all traders are Bayesians who correctly perceive the informativeness of the signals (*Bayesian pricing*). There are two separate forces that produce this overpricing. The first is simply belief heterogeneity: the highest valuation trader will be an over-reacter if there has been more good than bad news, and will be an under-reacter if the sequence of signals has more bad than good news. In either case, this highest valuation exceeds the valuation based on the correct Bayesian posterior on state A . The second source of overpricing is speculation. The equilibrium price will generally exceed the valuation of the most optimistic trader, because he/she believes in the possibility of a future sequence of public signals that would lead some other trader to be the most optimistic, at which point trade would occur and the currently most optimistic trader would cash out at a profit. We call the difference between the current equilibrium price and the maximum current valuation *the speculative premium*. The speculative premium is positive as long as it is still possible for the set of most optimistic traders to change at some future date.³

Another implication of our model concerns the trajectory of prices: *asymmetric reaction to good and bad news*. Because price responses are dampened when the marginal traders are under-reacters and exaggerated when the marginal traders are over-reacters, the difference between the price and 0.5

²Laboratory choice studies by economists and psychologists have consistently found a range of violations of Bayes's rule. For example, the study by El-Gamal and Grether (1995) classifies subjects into categories analogous to over-reacters and under-reacters.

³One can think of the speculative premium as representing a fair-odds bet by the currently most optimistic trader that he will profitably sell to a more optimistic trader at some later date.

is generally greater when there are more signals communicating good news versus bad news than when there is more bad news than good news.

To test the pricing predictions derived from our model, we run laboratory-controlled asset markets where asset returns are contingent upon a binary state of the world, and the information flows consist of a sequence of 10 informative public signals. In these markets, all traders are informed that the prior on state A is 0.5 and are provided with the conditional distribution of public signals given the state of the world. In one information treatment, the signals are highly informative, whereas signals are less informative in the other treatment. We impose the short-sale constraint and endow our traders with adequate liquidity, so liquidity constraints do not bind.

We find persistent and significant overpricing. That is, in both information treatments, we find pricing of the assets that is above the baseline of Bayesian updating to homogeneous posteriors. We also find that trading prices underreact to bad news compared to the reaction to good news, as implied by the model. We estimate a parametric model of the distribution of trader belief types, which allows us to test for heterogeneity of beliefs and also to back out estimates of the speculative premium. We find that the estimated speculative premiums are generally positive in those periods where the theory predicts it.

To dig more deeply into the overpricing phenomenon and to identify the extent to which it depends on the short-sale constraint, we run two additional variations on the simple one-market setup. In one variation, which we call *the complete markets treatment*, we open a second, complementary Arrow–Debreu security market that pays 1 in state B and 0 in state A . Traders are endowed with both assets and trading occurs simultaneously in both markets. Thus, good news for the A market is bad news for the B market, and vice versa. The choice of this alternative treatment has a number of motivations. First, the existence of overpricing is very easy to identify, because it is implied whenever the sum of the prices in the two markets exceeds 1. Second, past asset pricing experiments found that prices are sometimes distorted from rational expectations equilibrium prices when markets are incomplete, and very close to the rational expectations equilibrium when a complete set of Arrow–Debreu securities are traded (Plott and Sunder (1988)). Thus it is conceivable that the overpricing we observe in our baseline could have been due to market incompleteness. Third, for the same reason it is easy for the experimenter to identify overpricing when markets are complete, it is also easy for the traders to identify it. In particular, if prices add to more than 1, arbitrage opportunities become transparent, since any trader is able to sell one unit of each asset and make a sure profit, although these arbitrage opportunities are still limited by the short-sale constraint. We still find overpricing in these complete markets, which suggests that these kinds of possible effects cannot explain our consistent finding of overpricing in the baseline treatments.

This leaves the short-sale constraint as the most likely remaining explanation for the speculative overpricing we observe in our data. From a theoretical

standpoint, the short-sale constraint is an essential feature of the HK theory of speculative overpricing, and therefore a natural question is whether the overpricing persists if traders can sell the asset short. For this reason, in our final treatment, which we call *the short-sales treatment*, we continue to have both assets, but now allow short sales by permitting traders to buy from the experimenter unlimited units of a risk-free “bundled” asset, consisting of one unit of the A asset and one unit of the B asset, for a price of 1. To keep trading as simple a task as possible, we only allow market transactions in the A asset. However, this means that if the price of the A asset is higher than a trader’s valuation, that trader can buy a risk-free asset bundle and then sell the A unit of the bundle (retaining the B unit), generating an expected profit. Thus, this treatment relaxes the short-sale constraint, and indeed we observe traders buying the bundled asset and then unbundling it by selling the A asset portion; they are effectively shorting in the A market. This turns out to have a large downward effect on prices. We find lower prices in the short-sales treatment that are significantly closer to the homogeneous-belief Bayesian pricing than in either the baseline or the complete markets treatment.

Section 2 gives some background and discusses some of the related literature. The model and the theoretical results are presented in Section 3. Section 4 describes the experimental design and procedures. Results are presented in Section 5. Section 6 concludes with a summary of findings and suggestions for future work.

2. BACKGROUND AND DISCUSSION OF RELATED LITERATURE

2.1. *Asset Pricing Experiments*

There are three relevant classes of asset pricing experiments that provide a useful background and contrast with the experiment presented in this paper. First, there are a number of published multiperiod asset experiments that were designed to test rational expectations equilibrium with no uncertainty, where the asset paid off certain dividends in each period and perfect foresight pricing was easily calculated. These date back to the initial study by Forsythe, Palfrey, and Plott (1982; henceforth FPP). There is a connection with this paper, in that the pricing was determined by a very simple recursive calculation starting from the last period, and equilibrium had the property that, in each period, the price was determined by exactly one trader type who values the asset the highest. Forsythe et al. made two key findings in that experiment, which have been successfully replicated with a number of variations (Forsythe, Palfrey, and Plott (1984), Friedman, Harrison, and Salmon (1984)). First, prices converged over time toward the rational expectations prices. Second, prices always converged from *below*; that is, prices never exceeded the rational expectations prices. No speculative premium was ever observed. The current experiment differs from these earlier experiments by introducing state uncertainty, sequential public

information signals, and Arrow–Debreu securities that pay off only in the last period.

A second class of asset pricing experiments, initiated by Plott and Sunder (1982, 1988) and reviewed in Sunder (1995), explicitly focuses on the questions of whether and under what conditions state-contingent claims markets successfully aggregate private information in *static* markets, that is, rational expectations equilibrium in the sense of Radner (1979) and Grossman and Stiglitz (1980). Traders are endowed with private information at time 0, the market opens and clears at time 1, and, in theory, private information is fully revealed by the equilibrium price as if it had been public information from the start. Those experiments focus on questions about aggregation of private information and the conditions under which transaction prices converge to the fully revealing rational expectations equilibrium. One of their findings, which partly motivated our complete markets treatment, is that pricing was more consistent with the rational expectations theory when markets are complete, in the sense of including a full set of Arrow–Debreu securities, than when markets are incomplete as in our baseline treatments. More recent studies have dug deeper into questions about why standard predictions about price response to information (Asparouhova, Bossaerts, Eguia, and Zame (2009)) and the distribution of asset holdings (Bossaerts, Plott, and Zame (2007)) may fail. In contrast to the present paper, these approaches are based on the standard capital asset pricing model and explore the role of heterogeneity in attitudes toward risk and ambiguity, while our approach centers around heterogeneous beliefs.

The third class of experiments are the “bubble experiments” initiated by Smith, Suchanek, and Williams (1988). Like the first class, these are multi-period asset markets where the assets generate a stream of dividends. The dividends in each period are i.i.d. draws from a known distribution. Thus, unlike our model, realizations of the outcomes in each period provide no information about the future value of the asset. Rather, the expected value of the asset is known at all points in time, so there is no possibility for heterogeneous beliefs. Since dividends accrue each period, the fundamental asset value declines over time. Consequently, the equilibrium price dynamics for such markets are completely different from markets that share the properties of our model. In fact, if all traders are risk-neutral, equilibrium prices simply decline linearly to zero over time. If there are T periods remaining, the asset’s value is simply equal to T times the expected per-period dividend of the asset.

Indeed, the observed price dynamics in these bubble experiments are completely different from the equilibrium price dynamics in our model. The pricing more closely resembles the original FPP experiment. In early periods, transaction prices are *significantly below* the equilibrium price, as if there is a negative speculative premium. Because the equilibrium price declines over time while the price adjustment process drives the below-equilibrium prices upward, the transaction prices eventually catch up with equilibrium prices. When that happens, the price adjustment stops, and levels out. However, the *equilibrium* price

continues to fall. This results in a situation where prices exceed fundamental value—a bubble. The surprising observation in these experiments is that transaction prices often remain approximately constant for a while even though the fundamental value is declining. Volume declines as well, and then the price collapses to its fundamental value at or near the time the terminal period when the asset expires. This is obviously not an equilibrium phenomenon, at least within the class of models that motivated those experiments or the class of models considered here. A second finding from those experiments that mirrors the FPP class of experiments is that the disequilibrium pricing (both the *underpricing in early periods* and the *overpricing in middle-to-later periods*) diminishes with experience, leading to convergence in the direction of the rational expectations equilibrium. Also noteworthy is that equilibrium pricing in the basic bubble experiment does not depend on factors such as short sales or liquidity constraints, trader heterogeneity, complete markets, and so forth. In fact, researchers have run many variations, including futures markets and other types of market organization, which generally lead to similar conclusions. In one variation particularly relevant to the present paper (Porter and Smith (2003)), short sales are allowed, and the bubble phenomenon persists, and if anything is even more pronounced.

2.2. Theories of Speculative Trade in Asset Markets

Models in the finance literature have analyzed the impact of speculative trading due to heterogeneous beliefs on asset prices when no short sale is allowed. Biais and Bossaerts (1998) considered several types of heterogeneity in beliefs, such as common knowledge about the belief formation rules only, and derived the implied speculative value of the assets under each type. Scheinkman and Xiong (2003) found speculative bubbles with high volume and volatility in their model of differences in beliefs due to overconfidence. Our model is closest to the one studied in Harris and Raviv (1993), in which they looked at heterogeneity of beliefs in a model with a continuum of public signals, but where some traders have market power so prices are not determined competitively. Like Scheinkman and Xiong, they focused on the relationship between trading volume and price volatility.

Morris (1996) built a dynamic version of the HK speculative trading model to show that small differences in prior beliefs can lead to a significant speculative premium. In the HK model, the heterogeneity of expectations regarding others' beliefs that drove the speculative buying in anticipation of reselling was taken as given. Morris modeled this heterogeneity as initial differences in beliefs regarding the fundamental value of the asset, so that, as beliefs converge over time to the true probability, the speculative premium falls to zero as well. He also formalized Miller's (1977) claim that the most optimistic trader would hold all the assets, assuming sufficient liquidity, and that the most optimistic valuation would drive the equilibrium pricing. Ottaviani and Sorensen

(2007) analyzed the rational expectations equilibrium (REE) price dynamics in a binary prediction market where traders have heterogeneous priors and private information. They found that the prices actually under-react to information under the assumption that traders are liquidity-constrained or risk-averse. They also found that more information provided over time corrects this initial under-reaction so that the price approaches the Bayesian posterior. Finally, Asparouhova et al. (2009) explored the implications of a different kind of behavioral bias in beliefs by studying asset market equilibria with ambiguity-averse traders.

Our model builds on these ideas about speculation and belief heterogeneity and maintains the institutional assumptions of sufficient liquidity, risk-neutrality, and short-sale constraint. However, we depart from the assumption of heterogeneous priors and updating about the probability of future dividends based on the history of dividends, because ours is a model of an asset that pays off only at the end of the market based on the state of the world. Instead, traders in our model, who have a homogeneous prior, observe a sequence of public signals over the life of the asset, but draw different inferences about the state of the world from this information, which leads to heterogeneous posterior beliefs.

3. THE MODEL

Nature chooses the state of the world, $\varpi \in \{A, B\}$, where the probability of A is $p \in (0, 1)$. There is an asset market with $T + 1$ trading periods, $t = \{0, 1, 2, \dots, T\}$, and I risk-neutral traders, $i = \{1, 2, \dots, I\}$. There is one type of asset in this market. Each unit a trader holds at the end of period T pays off 1 if A is the state of the world and 0 if the state of the world is B . There are no intermediate direct returns from holding the asset in periods $0, \dots, T - 1$. Traders observe a sequence of public signals, $s = \{s_1, \dots, s_T\}$, where s_t denotes the signal observed at the beginning of trading period t . There are two sources of earnings in these markets: trading profits or losses from transactions made during the market and the one-time state-dependent payoff for the final asset holdings at the end of the market. Each trader is initially endowed with \bar{x}^i units of this risky asset and \bar{y}^i units of a safe asset that pays 1 in both states of the world (“cash”). We assume traders are risk-neutral, so if trader i 's final allocation of the risky asset is x_T^i , and final allocation of cash is y_T^i , then i 's utility is $E^i = y_T^i + x_T^i I_A$, where $I_A = 1$ in state A and $I_A = 0$ in state B .

Signals are binary, with $s_t \in \{a, b\}$, and are generated by a symmetric stochastic process that is independent and identically distributed across periods, conditional on the state.⁴ Conditional on $\varpi = A$, then $s_t = a$ with probability

⁴Most of the theoretical results hold for more general signal structures. Assumptions such as a binary signal space, independence, symmetry, and identical distributions over time are used for simplicity of exposition and to keep the theoretical model as close as possible to the experimental implementation.

$q > 0.5$ and $s_t = b$ with probability $1 - q$. Conditional on $\varpi = B$, $s_t = b$ with probability $q > 0.5$ and $s_t = a$ with probability $1 - q$. In the initial trading period, traders have no information about the state of the world except the prior p_0 . Since the asset pays off only in state A , we sometimes refer to the asset as *asset A* and sometimes refer to a signal $s_t = a$ as *good news* and a signal $s_t = b$ as *bad news*.

3.1. Equilibrium Prices With Bayesian Traders

First, suppose that all traders are Bayesians and use a common Bayesian updating rule, based on the “true” stochastic process generating the signals. That is, q is common knowledge and all traders update using Bayes’s rule. Let ρ_t be the common belief that the state of the world is A , given the history of signals $\{s_1, s_2, \dots, s_t\}$. Note that $\rho_0 = p$ because no information has yet been revealed. Given ρ_t , the common posterior at $t + 1$ if $s_{t+1} = a$ is

$$(1) \quad \rho_{t+1}^{s_{t+1}=a} = \frac{q\rho_t}{q\rho_t + (1-q)(1-\rho_t)},$$

and the common posterior at $t + 1$ if $s_t = b$ is

$$(2) \quad \rho_{t+1}^{s_{t+1}=b} = \frac{(1-q)\rho_t}{(1-q)\rho_t + q(1-\rho_t)}.$$

Given that the asset pays off 1 in state A and 0 in state B , and given that all agents are symmetric and risk-neutral, this common posterior is also the valuation of the asset. This is the Bayesian equilibrium price of the asset after any history.

3.2. Equilibrium With Heterogeneous Beliefs

This section contains a theory of pricing in the asset A market if traders have heterogeneous beliefs of a particular kind. As in the HK models, the traders agree to disagree. At every point in time, each trader thinks his own belief is absolutely correct. Traders have rational expectations about the distribution of future prices, in the sense that they agree on the mapping from sequences of signals to the equilibrium price, and disagree only about the fundamental value of the asset.

The traders could have *subjective priors* and start out with different home-grown prior beliefs p_0^i that the state is A . However, since we state clearly to the traders that states A and B are equally likely in the instructions, this type of belief heterogeneity is unlikely.

We focus on a model where different traders have different perceptions about the informativeness of each signal. In this case, traders start in period 0 with the common prior, p_0 , but each trader has his own personal estimate,

q^i , of the informativeness of the signal. These q^i 's could differ from the objective q of the signal.

This *subjective updating* leads to heterogeneity in the degree to which different traders will update their belief about the state of the world in response to identical sequences of signals. Specifically, it is possible that some traders over-react to news, and other traders under-react to news (relative to how a Bayesian with $q^i = q$ updates).⁵ Past experiments (e.g., Anderson and Sunder (1995), Goeree, Palfrey, Rogers, and McKelvey (2007), El-Gamal and Grether (1995)) have found evidence of this kind of judgment bias, including heterogeneity. Over-reaction to the signals is sometimes referred to as base-rate neglect or a base-rate fallacy, and under-reaction is sometimes referred to as conservatism (Camerer (2003)).

3.2.1. *Trader Types With Subjective Updating Heterogeneity*

Consider possible trader types characterized by the parameter $\theta \in [0, \infty]$. A trader with type θ_i will treat a single signal as if it had the informational equivalent of θ independent signals, each of informativeness q . Thus, θ_i measures how much trader i under-reacts ($\theta_i < 1$) or over-reacts ($\theta_i > 1$) to the signal relative to q .

Let ρ_{it} denote trader i 's belief at the beginning of period t that the state of the world is A given some history of public signals $\{s_1, s_2, \dots, s_t\}$. Since traders share a common prior when no information has yet been revealed, $p_{i0} = p$ for all $i \in I$. Given ρ_{it} , trader i 's updated posterior that the state of the world is A after observing $s_{t+1} = a$, if i is type θ_i , equals

$$(3) \quad \rho_{it+1}^{s_{t+1}=a}(\theta_i) = \frac{q^{\theta_i} \rho_{it}}{q^{\theta_i} \rho_{it} + (1 - q)^{\theta_i} (1 - \rho_{it})},$$

and after observing $s_{t+1} = b$ equals

$$(4) \quad \rho_{it+1}^{s_{t+1}=b}(\theta_i) = \frac{(1 - q)^{\theta_i} \rho_{it}}{(1 - q)^{\theta_i} \rho_{it} + q^{\theta_i} (1 - \rho_{it})}.$$

With this formulation of the trader types, posterior beliefs are always proper probabilities in the sense that trader i 's updated posterior that the state of the world is B after observing $s_{t+1} = a$ equals

$$(5) \quad 1 - \rho_{it+1}^{s_{t+1}=a}(\theta_i) = \frac{(1 - q)^{\theta_i} (1 - \rho_{it})}{q^{\theta_i} \rho_{it} + (1 - q)^{\theta_i} (1 - \rho_{it})},$$

and after observing $s_{t+1} = b$ equals

$$(6) \quad 1 - \rho_{it+1}^{s_{t+1}=b}(\theta_i) = \frac{q^{\theta_i} (1 - \rho_{it})}{(1 - q)^{\theta_i} \rho_{it} + q^{\theta_i} (1 - \rho_{it})}.$$

⁵This also implies that traders differ in their expectations about future signals.

We refer to traders with $0 \leq \theta < 1$ as *Skeptical* types. At one extreme is the $\theta = 0$ type. Traders of this type believe that the signals are just noise, as if the signal distribution were independent of the state. They do not update their prior after either signal a or signal b . Such a type's probabilistic belief that A is the state of the world remains unchanged for any sequence of signals. That is,

$$\rho_{it+1}^{s_{t+1}=a}(0) = p \quad \forall t, s_1, \dots, s_t.$$

We refer to traders with $\theta = 1$ as the *Bayesian* type. Traders of this type believe they are receiving signals of strength q , so their posterior probabilities are equivalent to those of a Bayesian:

$$\rho_{it+1}^{s_{t+1}=a}(1) = \frac{q\rho_{it}}{q\rho_{it} + (1-q)(1-\rho_{it})}.$$

We refer to $\theta > 1$ as *Gullible* types. Traders of this type update as if the informativeness of signals is higher than q . For extremely high values of θ , fickle traders treat a signal as nearly a full revelation of the state. For example, if $p = 0.5$, $q = 0.7$, and $\theta_i = 10$, then, after the first signal, if $s_1 = a$, trader i 's posterior is

$$\rho_{i1}^{s_1=a}(10) = \frac{0.7^{10}}{0.7^{10} + 0.3^{10}} = 0.9998.$$

This does not imply that gullible types' beliefs move immediately to 0 or 1 and stay there. In fact, exactly the opposite is the case. In the above example, if $s_2 = b$, then i 's beliefs go back to $\rho_{i2}^{s_1=a, s_2=b} = 0.5$, and then if $s_3 = b$ again, the trader's belief would be $\rho_{i3}^{s_1=a, s_2=b, s_3=b} = 0.0002$. Thus, gullible types have relatively volatile beliefs, while skeptical types have relatively sticky beliefs.

3.2.2. Equilibrium Prices

We maintain the assumptions of no short sale (implemented in the experimental design) and sufficient liquidity so that any trader can hold all units of the risky asset for any price less than or equal to 1. Under these assumptions, we can apply arguments similar to those used to solve the HK model and characterize the equilibrium price dynamics in our model.

For the remainder of the paper, we will assume $p = 0.5$.⁶ For any fixed p , the updating process depends only on q , the number of a signals, which we denote by α , and the number of b signals, which we denote by $\beta \equiv t - \alpha$. Hence, in the baseline case of homogeneous Bayesian beliefs ($\theta_i = 1 \forall i$), the equilibrium

⁶The model extends in a straightforward way to the more general case of $p \neq 0.5$.

price of the asset at period t , P_t^B , following any history in which the number of a signals is α , equals

$$P_t^B = \rho_t = \frac{q^\alpha(1 - q)^{t-\alpha}}{q^\alpha(1 - q)^{t-\alpha} + q^{t-\alpha}(1 - q)^\alpha}.$$

Given the way we have defined our different trader types, a trader's posterior beliefs will depend only on the trader's type, q , and the difference between the number of good news signals and bad news signals, $\delta = \alpha - \beta$. Specifically, the current belief of trader type θ_i can be expressed as

$$\begin{aligned} \rho_{it}^\alpha(\theta_i) &= \frac{q^{\theta_i\alpha}(1 - q)^{\theta_i(t-\alpha)}}{q^{\theta_i\alpha}(1 - q)^{\theta_i(t-\alpha)} + q^{\theta_i(t-\alpha)}(1 - q)^{\theta_i\alpha}} \\ &= \frac{1}{1 + \left(\frac{1 - q}{q}\right)^{\theta_i\delta}}. \end{aligned}$$

Define $\bar{\rho}_t(\alpha) = \max_{i \in \mathcal{I}} \{\rho_{it}^\alpha(\theta_i)\}$ to be the most optimistic belief among the traders in period t about A being the state of the world, and define $\theta_t^* = \arg \max_{i \in \mathcal{I}} \{\rho_{it}^\alpha(\theta_i)\}$. That is, $\bar{\rho}_t(\alpha) = \rho_{it}^\alpha(\theta_t^*)$. The equilibrium price of the asset at period t given the number of a signals, $P_t(\alpha)$, must be equal to the highest expected return of holding it to the next period. If the price is strictly lower than the highest expected return, then the trader(s) with the highest expected return would demand infinite units of the asset and the market would not clear. On the other hand, if the price is strictly higher than the highest expected return, then the demand for the asset would be zero and that price cannot be the equilibrium price.

Let $\varphi_t(\alpha)$ denote the most optimistic belief about the probability of an $s_{t+1} = a$ after α a signals, up to period t . Then,

$$(7) \quad \varphi_t(\alpha) = \bar{\rho}_t(\alpha) \left(\frac{q^{\theta_t^*}}{q^{\theta_t^*} + (1 - q)^{\theta_t^*}} \right) + (1 - \bar{\rho}_t(\alpha)) \left(\frac{(1 - q)^{\theta_t^*}}{q^{\theta_t^*} + (1 - q)^{\theta_t^*}} \right).$$

Note that this is not equivalent to the most optimistic belief about A being the state of the world, because $\varpi = A$ does not necessarily mean $s_{t+1} = a$. Traders can update their beliefs and asset valuations based only on the sequence of revealed signals, so pricing depends upon the revealed signals and the traders' expectations about future signals. The θ type with the most optimistic belief about the state of the world being A also has the most optimistic belief about the next signal being a .⁷ Now we can specify the equilibrium price

$$(8) \quad P_t(\alpha) = \varphi_t(\alpha)P_{t+1}(\alpha + 1) + (1 - \varphi_t(\alpha))P_{t+1}(\alpha).$$

⁷This follows because $\theta_t^* = \min_{i \in \mathcal{I}} \{\theta_i\}$ when $\bar{\rho}_t(\alpha) < 0.5$ and $\theta_t^* = \max_{i \in \mathcal{I}} \{\theta_i\}$ when $\bar{\rho}_t(\alpha) > 0.5$.

The first term on the right hand side (RHS) is equal to the most optimistic belief about an a signal being revealed next period multiplied by the price next period if $s_{t+1} = a$. The second term is equal to the corresponding belief that a b signal will be revealed next period multiplied by the price next period if $s_{t+1} = b$. Equation (8) states that the asset price must be equal to the highest expected return of holding the asset to the next period. If the price is lower than that, then the trader who would get the highest expected return would have infinite demand. If the price is higher than that, then no trader would want to hold the asset.

In the last period, period T , the price is equal to the highest posterior belief among all traders that the state is A :

$$(9) \quad P_T(\alpha) = \bar{\rho}_T(\alpha).$$

The equilibrium pricing scheme is uniquely pinned down by equations (6) and (7) because we can now solve backwards for the equilibrium price at every period. Note that our model and this specification of the equilibrium price dynamics depart from the original HK and Morris models in two specific ways. First, while they look at a finite truncation of an infinite market, we analyze a market with $T < \infty$ periods. Because of our finite horizon, we can immediately rule out other possible pricing trajectories involving bubbles or Ponzi schemes that Harrison and Kreps and Morris consider. Second, while their model introduces uncertainty as to whether the asset will pay a dividend after each period, the asset that we analyze pays off only at the end of the market after T periods. Thus, in their analysis, the price dynamics and speculative premiums are driven by heterogeneous beliefs about dividend payoffs in future periods based on the past dividend stream. In our analysis, the price dynamics and speculative premiums are driven by heterogeneous updating of beliefs about the state of the world that determines final asset payoff.

3.3. *Speculative Premium*

The above analysis shows that equilibrium prices in a market where traders have heterogeneous beliefs will typically be different from equilibrium prices if all traders are Bayesian ($\theta = 1$). The difference between equilibrium prices in our model and Bayesian pricing arises for two different reasons. At any point in time, trader k 's willingness to pay for the asset has two separate components: (a) a *valuation* component based on trader k 's current hold-to-maturity valuation; and (b) a *speculative premium* component that exists if there is some probability that trader k can resell the asset at some future date following a sequence of public signals that leads another trader to have a higher hold-to-maturity valuation than trader k .

Consider first how the valuation component affects prices when there are heterogeneous beliefs, in the sense that traders with different θ 's will have different hold-to-maturity valuations of the asset. Trader k 's valuation in period

t given α , $\rho_{kt}(\alpha)$, is simply his/her posterior belief that the state is A , which depends on θ_k . Any trader k is willing to pay *at least* his hold-to-maturity valuation, $\rho_{kt}(\alpha)$, and therefore the equilibrium price must be at least equal to $\bar{p}_t(\alpha) = \max_{i \in \mathcal{I}} \{\rho_{it}^\alpha(\theta_i)\}$. In states where $\alpha > \beta$, $\bar{p}_t(\alpha) = \rho_{it}^\alpha(\theta_{\max})$, and in states where $\alpha < \beta$, $\bar{p}_t(\alpha) = \rho_{it}^\alpha(\theta_{\min})$. Therefore, if $\theta < 1$ for some traders and $\theta > 1$ for other traders, as choice experiments have indicated, then this implies that equilibrium prices *must exceed Bayesian prices*, at least for all $\alpha \neq 0$.⁸

Consider next how the speculative premium component, the difference between the current equilibrium price and the current maximum hold-to-maturity valuation among all traders, affects prices when there are heterogeneous beliefs. That is, we define the *speculative premium* by $\pi_t(\alpha) = P_t(\alpha) - \bar{p}_t(\alpha)$. Following Morris (1996), we give two definitions of an optimistic trader (i.e., a trader with the maximum hold-to-maturity valuation).

DEFINITION 1: *Trader k is a current optimist at t if $\rho_{kt}(\alpha) = \bar{p}_t(\alpha)$.*

DEFINITION 2: *Trader k is a permanent optimist at t if $\rho_{kt'}(\alpha') = \bar{p}_{t'}(\alpha')$ for all $t' = t + 1, \dots, T$ and for all $\alpha' = \alpha, \alpha + 1, \dots, \alpha + t' - t$.*

In words, a permanent optimist at t not only has the (weakly) most optimistic belief among all traders at t that A is the state of the world, but will also continue to be an optimist for all possible sequences of future signals. The speculative premium can be calculated recursively by

$$\begin{aligned} \pi_t(\alpha) &\equiv \varphi_t(\alpha) [\pi_{t+1}(\alpha + 1) + \bar{p}_{t+1}(\alpha + 1)] \\ &\quad + (1 - \varphi_t(\alpha)) [\pi_{t+1}(\alpha) + \bar{p}_{t+1}(\alpha)] - \bar{p}_t(\alpha). \end{aligned}$$

It is straightforward to prove that $\pi_t(\alpha) \geq 0$ for all $t = 0, \dots, T$ and for all $\alpha = 0, 1, \dots, t$. The following result shows that the speculative premium is *strictly* positive if and only if there is no permanent optimist.

PROPOSITION 1: (i) *If $|\delta| < T - t$, then no trader is a permanent optimist and $P_t(\alpha) > \rho_{it}(\alpha) \forall i$, and $\pi_t(\alpha) > 0$.*

(ii) *If $|\delta| \geq T - t$, then there is a permanent optimist, and $\pi_t(\alpha) = 0$.*

See Appendix for the proof.

In our experimental setup, there are 10 signals released in each market, so $T = 10$. In this case, the condition for a positive speculative premium stated in part (i) of Proposition 1 simplifies to $\alpha < 5$ and $\beta < 5$. With fewer than five pieces each of both good and bad news, there is always the possibility of enough additional pieces of either good or bad news before the end of the

⁸In our model, $\rho_{kt}(\alpha) = 0.5$ for all traders if $\alpha = \beta$.

market such that the current optimist at period $\alpha + \beta = t$ will no longer be the optimist. However, if α is greater than or equal to 5, this is no longer possible. The θ_{\max} trader(s) is the permanent optimist because there will always be at least as many pieces of good news as there are pieces of bad news, regardless of future pieces of information. Similarly, if β is equal to or greater than 5, then the θ_{\min} trader(s) is the permanent optimist. The permanent optimist(s) will continue to hold the assets until the end of the market, so there is no speculative premium once a permanent optimist exists.

3.4. *Asymmetric Response to Good versus Bad News*

We also compare by how much the price at time t differs from the flat prior $p = 0.5$ when α pieces of good news have been revealed versus when $t - \alpha$ pieces of good news have been revealed. An implication of our model is that equilibrium prices react more to pieces of good news than pieces of bad news.

PROPOSITION 2: $1 - P_t(\alpha) < P_t(t - \alpha) \forall \alpha > \frac{1}{2}$.

See [Appendix](#) for the proof.

3.5. *Horizon Effect*

Next, we explore another pattern of the speculative premiums: the horizon effect. As the number of periods until the end of the market decreases, the speculative premium is nonincreasing. The first part of this horizon effect follows directly from Proposition 1: if a sufficiently large number of good or bad news signals have been revealed ($|\delta| \geq T - t$), then the speculative premium, $\pi_t(\alpha)$, will equal zero for all subsequent periods. This is true because with enough pieces of either good news or bad news, relative to the number of periods remaining, there is no possibility that the most optimistic trader will change, no matter how many pieces of good news or bad news follow.

The second part of the horizon effect is that, in periods where $|\delta| < T - t$, the speculative premium is nondecreasing in the horizon for fixed δ . With fewer trading periods left in the market, the probability of δ switching between positive and negative also decreases; therefore, the speculative premium cannot increase.

PROPOSITION 3: $\pi_t(\alpha) \leq \pi_{t-2}(\alpha - 1) \forall T \geq t > 1$ and $\alpha < t$.

See [Appendix](#) for the proof.

Note that since $t = \alpha + \beta$, the value of δ is the same at histories (α, t) and $(\alpha - 1, t - 2)$. Thus, the proposition states that the speculative premium is (weakly) higher in earlier periods, holding $\delta = \alpha - \beta$ constant.

4. EXPERIMENTAL DESIGN AND PROCEDURES

We began our experimental analysis by conducting six sessions of a one-asset trading market, which we refer to as the baseline sessions, with a total of 68 individual traders. The traders were registered Caltech students who were recruited by email solicitation. Sessions were conducted at the Social Science Experimental Laboratory at Caltech. Instructions were read out loud and screen displays were explained using a Powerpoint slide show in front of the laboratory at the beginning of each session. All interactions during the sessions took place through the computer interface. The trading interface used the open source software package Multistage Games.⁹

In each market of a session, a coin is flipped before the market opens to determine the state of the world: either State *A* (heads) or State *B* (tails). The result of the coin-flip is not announced until the market closes. We then organize and allow trading in a single-asset market, where each subject can take trading positions as buyers and/or sellers.¹⁰ To ensure adequate liquidity, all traders have a sufficiently large initial cash endowment. Traders are also endowed with three units of the asset. No short selling is allowed. There is also a bankruptcy constraint that does not allow any trader to engage in a transaction if her cash holdings go below zero. Each trader receives payoffs at the end of the market based on final asset holdings and cash holdings. All prices are in integers values. In state *A*, each unit of the asset pays off 100 experimental dollars at the end of the market; in State *B*, each unit of the asset pays off 0.¹¹

There are eleven trading periods in each market, each period lasting for 50 seconds. Trading is opened for the first trading period, and follows an open continuous double auction procedure. Subjects can type in bids to buy and/or offers to sell as many units of the asset as they want, subject to the liquidity and short-sale constraints. When a bid or offer is entered, it immediately shows up on the public bid and offer book, which is displayed in the center of each subject's screen. Only improving bids and offers could be made, and only the most recent current bid and offer are active. Subjects can accept a bid or offer by highlighting it with the mouse and clicking the "accept" button, subject to the bankruptcy and short-sales constraints. Subjects can also cancel an active bid or offer they had previously posted. At the 50-second mark, all unfilled bids and offers are cleared from the book, and the second 50-second trading period begins.

At the start of the second trading period, the binary public signal (good news or bad news) is drawn according to the distribution conditional on the

⁹<http://multistage.ssel.caltech.edu/>.

¹⁰Additional procedural details are available in the Supplemental Material (Palfrey and Wang (2012)).

¹¹In four of the baseline sessions, the state 2 payoffs equaled 20 instead of 0. In the analysis of data, all transactions and prices are rescaled on a 0 to 100 scale. Experimental dollars were converted to U.S. dollars using an exchange rate of either 0.01 or 0.02, depending on the session.

original coin-flip and publicly announced to all subjects. Holdings are carried over across periods. Trading occurs in the second period following the same rules and procedures as in the first period. After 50 seconds, the book is again cleared and a new public signal is drawn and announced. This continues for 11 trading periods (until 50 seconds after the 10th public signal has been announced). After the last trading period, the market closes, the state of the world is revealed, and each trader's cash on hand is credited based on final holdings of the asset. We then proceed to open another market, with procedures identical to the first market. The experimenter again flips a coin to determine the state, trading screens are reset, asset endowments are reset at three units for each trader, and cash holdings are carried over from the first market. This continues until a total of six markets are conducted. Each subject is then paid in private the sum of his or her earnings in all six markets, plus a show-up fee of \$10. Each session lasted between 1.5 and 2 hours, including instructions and payment.

The public signal was generated by rolling a die in each period, as described in the instructions (see [Appendix](#)). In three of the sessions, the signal distribution corresponded to an informativeness of $q = \frac{5}{9}$, and in the other three sessions, the signal informativeness was $q = \frac{6}{9}$. These conditional signal distributions were explained carefully and accurately to the subjects.

Two additional sessions with 10 traders in each were conducted using a complete market design. In these sessions, traders were allowed to buy and sell two assets, one which paid off in state A and the other in state B . Hence, these markets offered the opportunity for limited arbitrage, suggesting the hypothesis that speculative overpricing will be diminished compared to the baseline sessions. Traders were endowed with three units of each asset. In other respects, they were conducted in an identical manner as the baseline sessions described above.

To explore the effect of the short-sale constraint on asset prices directly, we conducted three additional sessions where markets were organized to allow traders to engage in short sales. Specifically, at any time the market was open, any trader was allowed to purchase from the "bank" a safe, bundled asset consisting of one unit of asset A and one unit of asset B at a risk-free price of 100. Traders were allowed to purchase as many units of the safe, bundled asset as they wished, subject to the cash-on-hand constraint.¹² This allows any trader who has zero asset A holdings to engage in a strategy that mimics short-selling asset A , by purchasing the bundled asset and then unbundling it by selling off the asset A part. All three sessions used the signal strength $q = \frac{5}{9}$ and there were six 11-period markets. The procedures were otherwise the same as the one-asset market sessions: traders could hold units of both A and B assets, but only the A market was open for trading. [Table I](#) provides a summary of the experimental design.

¹²Traders could also resell the bundled asset back to the bank for a price of 100.

TABLE I
SUMMARY OF EXPERIMENTAL SESSIONS

Session	Type	Signal	# Mkts	# Traders	Period (sec)
1	One asset	s s s s s s s s s s s	6	12	50
2	One asset		6	12	50
3	One asset		6	10	50
4	One asset		6	12	50
5	One asset		6	12	50
6	One asset		6	10	50
7	Complete		6	10	50
8	Complete		6	10	50
9	Short sale		6	11	50
10	Short sale		6	12	50
11	Short sale		6	12	50

5. RESULTS

We present the results of our experiment in the following order. First, we analyze whether there is significant overpricing in the data from the baseline sessions, and to what extent this can be attributed to a speculative premium. We next test two related implications of the model: asymmetric reaction to good news versus bad news, and the horizon effect. Third, we report results from our two alternative market designs, which allow inframarginal traders, who believe prices are too high relative to their beliefs, to engage in arbitrage in ways that were not available in the baseline design. This includes completing the market by opening up markets for state contingent claims in both states, and explicitly allowing short sales. Finally, we examine the dynamics of individual asset ownership from period to period as information is gradually revealed, and compare it to the theoretical predictions about ownership dynamics.

The hypotheses generated by the model concern the trajectory of transaction prices in the markets. The central hypotheses concern overpricing relative to the Bayesian benchmark, which, as we showed in Section 3, is driven by two different phenomena: the valuation component, which derives from heterogeneous *current* hold-to-maturity valuations, and a speculative component, which derives from heterogeneous *future* hold-to-maturity valuations. Because we cannot control for, or directly observe, each individual trader's θ_i , measuring these two components of overpricing in our data means that the distribution of θ_i 's has to be estimated from the observed transaction prices. Before proceeding with the estimation, we first ask simply whether, in our markets, the combined effect of the two components produces prices that are in excess of the prices that would arise if traders had homogeneous, Bayesian beliefs.

We then turn to speculative overpricing by estimating the speculative premium. To do this, we estimate θ_{\min} and θ_{\max} for each session. We also obtain estimates for a model of homogeneous beliefs by constraining $\theta_{\min} = \theta_{\max}$. This allows us to conduct a nested test to see whether our data reject the null hypothesis of homogeneous beliefs. This is done separately for each session. We then use these session-by-session estimates to obtain a quantitative measure of the implied speculative premium component of overpricing in every trading period of every session, that is, the difference between the price and the valuation of the most optimistic type. This provides a direct test for the existence of a speculative premium in our data, as well as allowing us to test the model's specific theoretical predictions about how the speculative premium depends on the history of public signals. We also provide some finer tests of the implications of the theoretical model based on heterogeneous beliefs.

5.1. *Transaction Prices*

5.1.1. *Overpricing Relative to the Bayesian Benchmark*

To begin our analysis, we compare asset prices to the Bayesian benchmark, that is, the value of the asset assuming Bayesian updating ($\theta = 1$). We calculate the median price of all transactions in each trading period and use this as our price observation for that trading period. For the analysis in this section of the paper, we relate prices by the amount of information revealed. To do this, we code the history of public signals that has been revealed up to period t by counting the number, α , of good news signals and the number, β , of bad news signals. The observations for our analysis are aggregated at the period level. However, for ease of presentation in this section of the analysis, we construct an aggregate price for all periods in all markets that share the same δ . That is, we use the median of the median transaction prices over all trading periods with the same value of δ . The $\delta = 0$ trading periods are further broken down into two categories, depending on whether it was the initial trading period of a market ($\alpha = \beta = 0$) or a later trading period ($\alpha = \beta > 0$).

Signal Strength: $q = \frac{5}{9}$. Table II presents the aggregate median prices for the $q = \frac{5}{9}$ sessions, for each value of δ as well as the predicted prices for the homogeneous Bayesian updating model. N refers to the total number of transactions that occurred at that value of δ .

Figure 1 plots the prices in the $q = \frac{5}{9}$ sessions against the difference in good versus bad news signals, along with the predicted prices under the null model where all traders Bayesian update to the common posterior after receiving each signal (i.e., no heterogeneity). The observed transaction prices remain above the predicted ones regardless of the difference between good and bad news signals. Furthermore, the observed and predicted prices under the null model are significantly different from each other, according to the Wilcoxon

TABLE II
 MEDIAN PRICES BY INFORMATION REVEALED ($q = \frac{5}{9}$)

δ	Median Price (N)	Bayesian Price ($\theta = 1$)
-6	37.5 (4)	20.8
-5	48.1 (4)	24.7
-4	51.2 (24)	29.1
-3	50 (52)	33.9
-2	52.5 (113)	39
-1	55.6 (182)	44.4
0 (initial period)	53.8 (136)	50
0	58.8 (148)	50
1	60 (138)	55.6
2	61.2 (58)	61
3	75.2 (24)	66.1
4	87.5 (35)	70.9
5	93.1 (7)	75.3

signed-rank test ($p < 0.0001$).¹³ This suggests that although the traders are receiving informative signals about the state of the world, they may be using non-

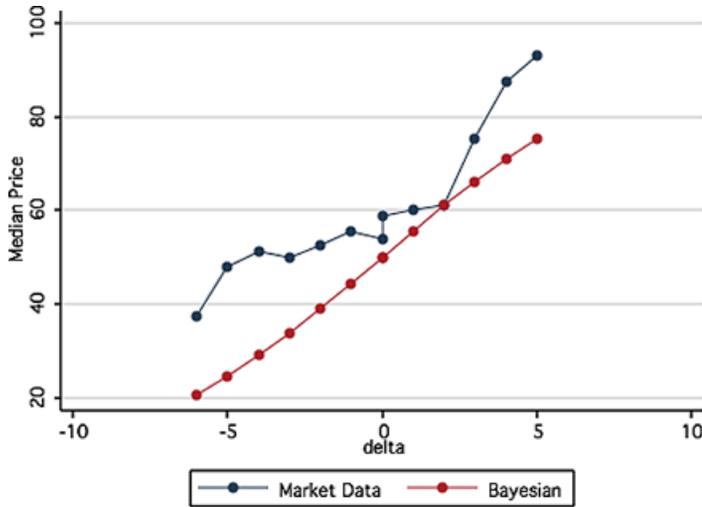


FIGURE 1.—Median prices versus Bayesian predictions ($q = \frac{5}{9}$).

¹³The Wilcoxon test assumes independence across observations. To the extent that there is some correlation across observations in our data, the p -values we report for these tests should be treated as a lower bound. Alternative tests with lower power (e.g., binomial test) also yield highly significant test statistics.

Bayesian updating heuristics. To the extent that there is heterogeneity of these heuristics, asset prices will deviate systematically from those predicted under the assumption of perfect Bayesian updating and lead to overpricing according to the multiple- θ model of speculation.

The homogeneous beliefs model predicts the price to be 50 in all periods where there are equal pieces of good and bad news, $\delta = 0$. A Wilcoxon signed-rank test reveals that the median prices in these periods are significantly higher than 50 ($p < 0.0001$). Next, we turn to the price in the initial period when there have been no news announcements. Under the null model, the price in the initial period of each market should be 50, reflecting the flat prior. This prediction does not hinge upon any assumptions about the belief updating process. In fact, the median price is above 50 in all 18 initial periods (Wilcoxon signed-rank: $p = 0.0002$). These transaction prices in the initial periods may offer the clearest evidence of speculative trading. Since no information has been revealed, if the prices are above 50, at least some traders must be trading based on speculation about price changes in future periods.

Signal Strength: $q = \frac{6}{9}$. Table III and Figure 2 display the aggregate median prices for the $q = \frac{6}{9}$ sessions. We find that observed prices follow the trajectory of predicted Bayesian prices more closely than in the $q = \frac{5}{9}$ sessions. However, the observed prices are still greater than the predicted prices for nearly all δ , and these differences are significant according to the Wilcoxon signed-rank test ($p < 0.0001$).

TABLE III
MEDIAN PRICES BY INFORMATION REVEALED ($q = \frac{6}{9}$)

δ	Median Price (N)	Bayesian Price ($\theta = 1$)
-7	0.9 (6)	0.8
-6	2.2 (8)	1.5
-5	4.4 (2)	3
-4	10.2 (12)	5.9
-3	13.1 (30)	11.1
-2	21.9 (57)	20
-1	43.8 (80)	33.3
0 (initial period)	58.1 (110)	50
0	62.5 (98)	50
1	75 (146)	66.6
2	87.5 (88)	80
3	92.5 (53)	88.9
4	97.5 (39)	94.1
5	98.8 (5)	97
6	98.1 (8)	98.5
7	99.4 (2)	99.2
8	100 (1)	99.6

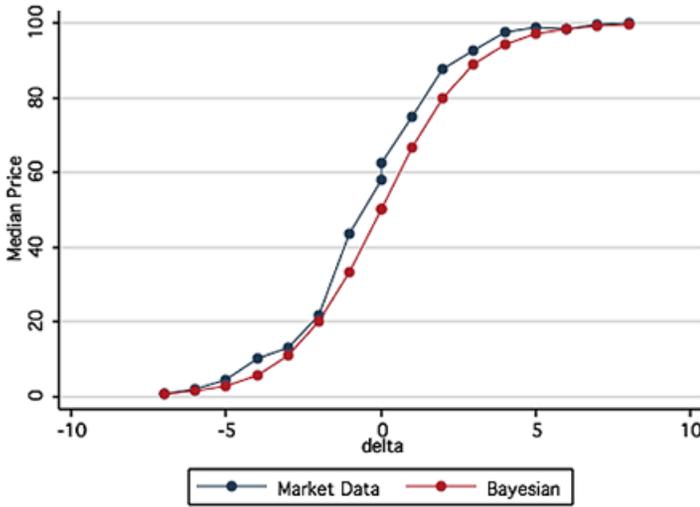


FIGURE 2.—Median prices versus Bayesian predictions ($q = \frac{5}{9}$).

We again look at the median prices in all periods where $\delta = 0$ and the prices are predicted to be 50 under the null model with homogeneous Bayesian updating, just as in the $q = \frac{5}{9}$ treatment. Here again we find that the median prices are also significantly higher than 50 (Wilcoxon signed-rank: $p < 0.0001$). The median prices in the 18 initial periods are not as uniformly high as they were in the $q = \frac{5}{9}$ treatment. Nevertheless, they are still significantly higher than 50 (Wilcoxon signed-rank: $p = 0.0120$), with only four periods having a median price lower than 50 and three more periods exactly at 50.

RESULT 1: Prices in the one-asset market are systematically higher in all treatments than equilibrium prices based on the null model of correct homogeneous trader beliefs about q .¹⁴

5.1.2. Speculative Overpricing

In this section, we focus on a key implication of our model, speculative overpricing. We estimate θ_{\min} and θ_{\max} from the pricing data for each session, which we then use to obtain quantitative measures of the speculative premium as a function of the history of public signals. We test the implications of Propositions 1 and 3 on the estimated speculative premiums.

¹⁴The transaction prices are also inconsistent with a model of homogeneous *but incorrect* beliefs about q (i.e., homogeneous θ , but $\theta \neq 1$). For the model with homogeneous $\theta < 1$, prices are predicted to be *less than* the null model when $\delta > 0$; and for the model with homogeneous $\theta > 1$, prices are predicted to be *less than* the null model when $\delta < 0$. Both predictions are rejected in our data.

Estimating the Distribution of θ -Types. We use our transaction price data to estimate the maximum and minimum θ types¹⁵ for each session using the following procedure. Using the recursive formulae in our model described in Section 3, we can calculate the theoretical price trajectory given any pair of values $(\theta_{\min}, \theta_{\max})$ for every sequence of signals in our data. Depending on the sequence of public signals, either the maximum θ type or the minimum θ type will be the most optimistic traders, and this is what determines the asset price trajectories in our model. Recall that $\theta = 0$ corresponds to a trader who acts as if signals contain no useful information about the state, and $\theta = 1$ corresponds to a trader who Bayesian updates with the correct q . We now compute the equilibrium price trajectories for all pairs of $\theta_{\min} = 0, 0.1, 0.2, \dots$ and $\theta_{\max} = 0, 0.1, 0.2, \dots$ such that $\theta_{\min} \leq \theta_{\max}$. This produces a matrix of prices that depends on α and β . Note that our estimation procedure also allows for the constrained model of homogeneous beliefs where $\theta_{\min} = \theta_{\max}$. This implies a nested test for heterogeneous beliefs. With homogeneous beliefs, there is no speculative premium.

For each possible $(\theta_{\min}, \theta_{\max})$ pair, we sum up the squared deviations of the median price in each trading period of each market from the theoretical price for that pair. Formally, let P_{gmt} be the median transacted price in trading period t of market m of session g . Let α_{gmt} and β_{gmt} denote, respectively, the number of a signals and b signals received up to and including period t in market m of session g . Let $P_t^*(\alpha_{gmt}, \beta_{gmt} | \theta_{\min}, \theta_{\max})$ denote the equilibrium prices from our theoretical model. Then we define the model error as the sum of squared deviations of the price data in session g from the theoretical model, evaluated at parameters $(\theta_{\min}, \theta_{\max})$:

$$e_g(\theta_{\min}, \theta_{\max}) = \sum_{m,t} [P_{gmt} - P_t^*(\alpha_{gmt}, \beta_{gmt} | \theta_{\min}, \theta_{\max})]^2.$$

The estimated parameters of the model for session g are given by

$$(\hat{\theta}_{\min}^g, \hat{\theta}_{\max}^g) = \arg \min_{0 \leq \theta_{\min} \leq \theta_{\max}} \{e_g(\theta_{\min}, \theta_{\max})\}.$$

Both the observed and predicted price are normalized to $[0, 1]$. We also pooled the sessions of each treatment together to estimate a treatment-level $\hat{\theta}_{\min}$ and $\hat{\theta}_{\max}$. The results are displayed in Table IV.¹⁶

¹⁵The theory does not depend on the distribution of types except for the minimum and the maximum θ .

¹⁶In 12 out of 396 trading periods, there were no transactions. These are treated as missing data.

TABLE IV
TYPE ESTIMATION BY SESSION^a

Session	q	$\hat{\theta}_{\min}^g, \hat{\theta}_{\max}^g$	$\hat{\theta}_{\min}^g = \hat{\theta}_{\max}^g$	F -stat
1	5/9	0.2, 1.7	0.6	236.1*
2	5/9	0, 1	0.3	123.1*
3	5/9	0, 1.8	0.3	227.3*
Pooled	5/9	0, 1.5	0.5	484.4*
4	6/9	0.5, 1	0.9	46.6*
5	6/9	1.1, 1.3	1.1	1.0
6	6/9	0, 1.5	0.8	131*
Pooled	6/9	0.6, 1.3	0.9	78.6*

^a * indicates homogeneous model (no speculative premium) rejected ($p < 0.05$).

Column 3 of Table IV shows the estimated $(\hat{\theta}_{\min}^g, \hat{\theta}_{\max}^g)$ pairs. Column 4 displays the best fitting homogeneous θ model. Column 5 contains the F -test statistic for the null hypothesis $\hat{\theta}_{\min}^g = \hat{\theta}_{\max}^g$, where

$$F = \frac{e_g(\theta_{\min} = \theta_{\max}) - e_g(\theta_{\min} < \theta_{\max})}{(n - 1) - (n - 2)} \cdot \frac{e_g(\theta_{\min} < \theta_{\max})}{n - 2},$$

and n is the number of trading periods. The fit is uniformly worse for the single θ estimations compared to the θ pair ones, suggesting that our model with heterogeneous posterior beliefs among the traders better captures the observed price dynamics. The homogeneous belief model is rejected at the 5% level for all sessions except one. In the one exception, the price dynamics suggest that most of the traders' beliefs are close to the objective signal strength.

Some additional observations can be gleaned from the estimation results for both the heterogeneous and homogeneous belief models. For the homogeneous belief model, $\hat{\theta}_{\max}^i = \hat{\theta}_{\min}^i$ is less than 1 for five of the six sessions. If we assumed that all traders updated their beliefs in the same way, then the price trajectory would suggest that, on average, traders under-reacted to the information flow. On the other hand, under the heterogeneous belief model, $\hat{\theta}_{\max}^i$ is estimated to be at least 1 for all six sessions. Furthermore, $\hat{\theta}_{\min}^i$ and $\hat{\theta}_{\max}^i$ span a considerable range for most of the sessions. Given the better fit of this model for nearly all the sessions, the likely heterogeneity across traders, ranging from those who react little to information to those who over-react, would have been masked by a homogeneous belief model.¹⁷

¹⁷We would expect such variation within and across sessions if there is some underlying distribution of θ^i in the population and we are drawing small (10 or 12) independent samples of traders from this distribution.

RESULT 2—Heterogeneity: We estimate significant heterogeneity in updating rules across subjects in all sessions. The range varies across treatments, with $\hat{\theta}_{\max}^i$ greater than or equal to 1.

To illustrate how the observed sequences of transacted prices compare to the prices in the estimated model with heterogeneity, Figure 3 displays price graphs of five asset markets from four different sessions. The markets were chosen to illustrate a wide range of behaviors and a range of different es-

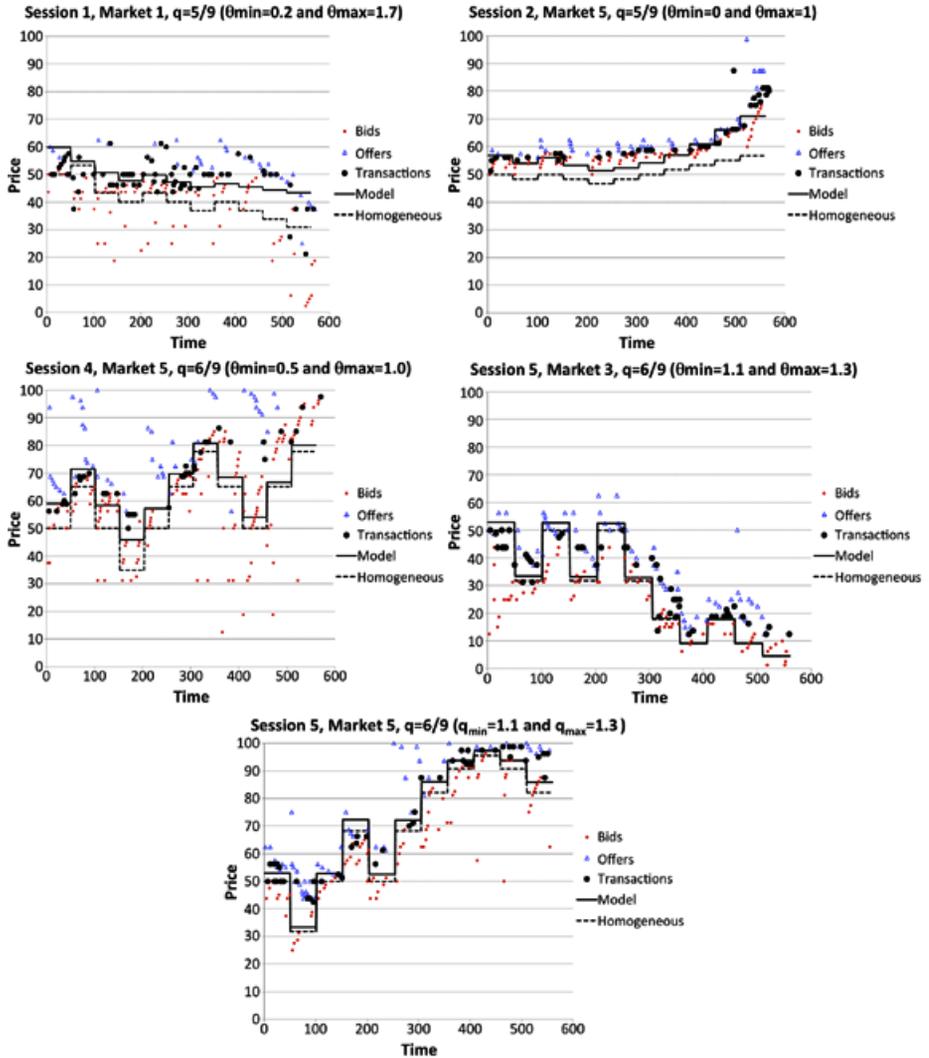


FIGURE 3.—Five sample markets from four sessions.

timated θ 's. We include two markets from the same session to illustrate the session-level model fit on markets with very different information flows. These graphs present all the bids, offers, and transactions in each market, in addition to the price trajectory of our model estimates as well as of the best fitting homogeneous θ for that session. Transacted prices appear as large dots in the graph, unaccepted bids to buy appear as small dots, and unaccepted offers to sell appear as small triangles. The estimated prices from our model appear as solid lines and the estimate prices from the best fitting homogeneous θ model appear as dashed lines.

Estimating the Speculative Premium. We use the session-specific $\hat{\theta}_{\min}^s$ and $\hat{\theta}_{\max}^s$ to calculate the speculative premium for each period of each session. Recall that the speculative premium is the difference between the price and the maximum valuation of the asset among all traders, which is determined by either the θ_{\min} or θ_{\max} trader, depending on the information revealed: $P_i(\alpha) - \bar{p}_i(\alpha)$. For each period, we determine whether the θ_{\min} traders or the θ_{\max} traders are the marginal ones. We then calculate the maximum valuation of the asset by these marginal traders given their θ , and subtract that from the period's median asset price. This difference is the trading period's speculative premium. We then take the median of all speculative premiums in periods with the same δ . Table V presents the median speculative premium as a function of δ , the analogue to the price Tables II and III.

We can see from Table V that the speculative premium is higher in periods with δ closer to 0, with a few exceptions. This overall pattern is consistent with

TABLE V
SPECULATIVE PREMIUM BY INFORMATION REVEALED

δ	$q = \frac{5}{9} (N)$	$q = \frac{6}{9} (N)$
-7		-7.18 (6)
-6	-5.85 (4)	-8.92 (8)
-5	3.68 (4)	-10.65 (2)
-4	5.00 (24)	-0.70 (12)
-3	3.34 (52)	1.09 (30)
-2	4.73 (113)	1.81 (57)
-1	5.63 (182)	10.45 (80)
0 (initial period)	3.75 (136)	8.13 (110)
0	8.75 (148)	12.50 (98)
1	2.59 (138)	3.61 (146)
2	0.27 (58)	1.11 (88)
3	-0.16 (24)	0.028 (53)
4	7.81 (35)	0.97 (39)
5	6.17 (7)	1.78 (5)
6		-0.34 (8)
7		0.15 (2)
8		0.39 (1)

Proposition 1(i): speculative overpricing should be reflected in a *positive speculative premium* in periods where $\alpha < 5$ and $\beta < 5$ since there are no permanent optimists. The median speculative premium for these periods is 5.00 in the $q = \frac{5}{9}$ treatment, 3.38 in the $q = \frac{6}{9}$ treatment, and 5.00 overall, all significantly positive (Wilcoxon signed-rank test: $p < 0.05$). The speculative premiums are also significantly positive across all trading periods (Wilcoxon signed-rank test: $p < 0.05$) in the $q = \frac{5}{9}$ treatment (median: 4.80), in the $q = \frac{6}{9}$ treatment (median: 1.89), and pooled across both treatments (median: 3.75).

RESULT 3—Proposition 1(i): The speculative premium is significantly positive in periods with no permanent optimist.

The second part of Proposition 1 states that the speculative premium should be zero when enough Good News signals (or Bad News signals) have accumulated relative to the number of periods remaining in the market ($T - t \leq |\delta|$). We first test a weak form of this hypothesis, that is, the speculative premium should be higher when $T - t > |\delta|$ than in periods where the horizon is too short ($T - t \leq |\delta|$). Indeed, this is what we observe: the speculative premiums are higher, on average, in periods where both α and β are less than $\frac{T}{2} = 5$ (Table VI), and this difference is significant in both treatments. Second, a somewhat stronger prediction is that the speculative premium should be positive if and only if both $\alpha < 5$ and $\beta < 5$. Across all periods where either α or β is greater than $\frac{T}{2} = 5$ ($T - t \leq |\delta|$), the speculative premiums are only significantly different from zero in the $q = \frac{5}{9}$ treatment. The null hypothesis that the speculative premiums are zero cannot be rejected for the $q = \frac{6}{9}$ treatment or when we pool across both information treatments. Thus, with the one exception of trading periods when $T - t \leq |\delta|$ in the $q = \frac{5}{9}$ treatment, we find strong support for our hypothesis.

RESULT 4—Proposition 1(ii): The speculative premium is not significantly different from zero in periods where either α or β is greater than $\frac{T}{2} = 5$ ($T - t \leq |\delta|$), with the exception of periods where $T - t \leq |\delta|$ in the $q = \frac{5}{9}$ treatment.

TABLE VI
MEDIAN SPECULATIVE PREMIUMS^a

	$q = \frac{5}{9}^{\dagger}$	$q = \frac{6}{9}^{\dagger}$	Pooled [†]
$\alpha \geq 5$ or $\beta \geq 5$	3.19*	0.99	2.00
$\alpha < 5$ and $\beta < 5$	5.00*	3.38*	5.00*

^a * = significantly different from 0 ($p = 0.05$). [†] = significantly lower in periods with $\alpha \geq 5$ or $\beta \geq 5$ ($p = 0.05$).

5.1.3. *Asymmetric Pricing in Good News versus Bad News Regimes*

The asymmetric price response to information in good news regimes compared to bad news regimes is already evident in Table II and Figure 1. The median price never goes below 30 for information flows in bad news regimes ($\delta < 0$), while the price reaches above 90 in good news regimes ($\delta \geq 0$).

The intuition behind our pricing asymmetry hypothesis (Proposition 2) is that the marginal trader is a high- θ type in good news regimes and a low- θ type in bad news regimes. Thus, prices react more strongly in good news regimes (δ positive) than in bad news regimes (δ negative). If traders had homogeneous beliefs, even if the common θ were not 1, the prices would be above the Bayesian benchmark only in good news regimes (if $\theta > 1$) or only in bad news regimes (if $\theta < 1$), but not in both.

To test Proposition 2 more carefully, we run regressions to test if the price is indeed less sensitive to bad news than to good news. The dependent variable is the deviation of the median price from 50 in each period. This is calculated by subtracting the price from 50 if $\delta < 0$ and subtracting 50 from the price if $\delta \geq 0$.¹⁸ The independent variables are interaction terms, one between the absolute difference between a signals and b signals, $|\delta|$, and a dummy for this difference being negative, and another between the difference and a dummy for a nonnegative difference. We run the regression separately for each treatment and again with the treatments pooled:

$$|50 - P| = \pi + \gamma_1 |\delta| * I(\delta \geq 0) + \gamma_2 |\delta| * I(\delta < 0) + \epsilon.$$

We hypothesize that $0 \leq \gamma_2 < \gamma_1$ because, from Proposition 2, the price should be further from 50 if $\delta \geq 0$ than if $\delta < 0$. Table VII reports the regression results.

We find that γ_2 is significantly less than γ_1 in both the separate and pooled regressions as hypothesized ($p < 0.05$). When there are already at least as many a signals as b signals ($\delta \geq 0$), the estimated price reaction to an additional a signal ($\hat{\gamma}_1$) ranges from 7.35 to 9.40, and is significantly greater than zero in all

TABLE VII
PRICE REACTION TO GOOD NEWS VERSUS BAD NEWS^a

	5/9**	6/9*	Pooled**
γ_1	7.35 (0.74)	8.40 (1.02)	9.40 (0.56)
γ_2	-2.53 (1.29)	5.45 (1.18)	1.93 (0.61)
Constant	3.10 (1.08)	11.66 (2.19)	5.85 (1.03)

^a $\gamma_1 > \gamma_2$ (*: $p = 0.05$; **: $p = 0.01$). Clustering by market.

¹⁸We run the same regressions with the absolute deviation from 50 as the dependent variable and the qualitative results remain the same.

cases ($p < 0.01$). In contrast, if there are already fewer a signals than b signals ($\delta < 0$), the estimated price reaction to an additional b signal ($\hat{\gamma}_2$) is not significantly different from zero in the $q = \frac{5}{9}$ treatment and is significantly greater than 0 for the $q = \frac{6}{9}$ treatment and for the pooled sample. We also note that the constant is significantly greater than 0 in all the regressions, which indicates speculative overpricing when $\delta = 0$.

Finally, we test one additional implication of the asymmetric response hypothesis: we compare the absolute price change when δ increases from 1 to 2 versus when δ decreases from -1 to -2 .¹⁹ Pooled across both treatments, the absolute price change when δ changes from 1 to 2 is 8.38 on average ($N = 27$). This is significantly higher than the absolute price change when δ changes from -1 to -2 , which is 5.54 on average ($N = 31$; Mann–Whitney: $p < 0.05$).

RESULT 5—Proposition 2: Market prices react asymmetrically to information in good and bad news regimes, and this asymmetry is consistent with the equilibrium price dynamics predicted by the heterogeneous θ updating model.

5.1.4. *The Horizon Effect*

The horizon effect posits that the speculative premium is larger when there are more periods remaining in the market, provided there is not yet a permanent optimist. Formally, when $T - t > |\delta|$ (i.e., for our markets, if $\alpha < 5$ and $\beta < 5$), the speculative premium is weakly increasing in $T - t$ for any fixed value of δ .

To test Proposition 3, we first construct a *horizon* measure, which we specify as the number of trading periods that remain in the market; thus, it ranges from 10 for the initial period to 0 for the last period. For each treatment, we regress the estimated speculative premium on the horizon variable, controlling for the difference in the pieces of good versus bad news, δ . We report the results separately for $\delta \leq 0$ and $\delta \geq 0$, because of the asymmetry of prices depending on whether δ is greater or less than zero. These regressions were restricted to the periods where $\alpha < 5$ and $\beta < 5$ because our theory only predicts the horizon effect for these periods. The regression coefficients are reported in Table VIII.

According to Proposition 3, the coefficient on $T - t$ in the regression should be greater than or equal to zero. We find that three of the four coefficients on the horizon variable are not significantly different from zero, and the one significant coefficient has the opposite sign, as predicted by Proposition 3. The coefficients on $|\delta|$ are all significant and negative, indicating that the speculative premium is smaller the further the posterior belief is from 0.5. This seems intuitive, since the speculative premium arises because of the possibility that the identity of the most optimistic trader may switch at some future date. The

¹⁹There are too few observations to compare $\delta = 2$ to $\delta = 3$ and $\delta = -2$ to $\delta = -3$ (or higher levels of δ).

TABLE VIII
THE HORIZON EFFECT IN SPECULATIVE PREMIUMS FOR PERIODS
WHERE $\delta \geq 0$, $\alpha < 5$, AND $\beta < 5^a$

	$\delta \geq 0$		$\delta \leq 0$	
	$q = \frac{5}{9}$	$q = \frac{6}{9}$	$q = \frac{5}{9}$	$q = \frac{6}{9}$
$T - t$	-0.45 (0.23)	-1.22* (0.51)	-0.42 (0.33)	-0.066 (1.18)
$ \delta $	-4.06** (0.58)	-5.13** (1.44)	-1.47* (0.67)	-5.00* (2.11)
Constant	10.54** (1.98)	18.60** (5.65)	9.99** (2.83)	16.75 (12.31)
N	84	96	108	82

^aSignificantly different from 0 (*: $p < 0.05$; **: $p < 0.01$). Clustering by market.

probability of such a switch is decreasing in δ if $\delta > 0$ and increasing in δ if $\delta < 0$. Furthermore, the switch, if it occurs, would happen closer to the horizon the larger is $|\delta|$, controlling for $T - t$.

RESULT 6—Proposition 3: We find no significant horizon effect, except in one case where the speculative premium is significantly increasing in t .

5.2. Complete Markets and Relaxing the Short-Sale Constraint

5.2.1. Complete Markets: Both Assets Traded

We compare the price trajectories in the two markets in the complete markets environment to those in the incomplete market environment where only one asset is traded. The prices in the complete markets do reach substantially lower levels, in the 20s and 30s when $\delta < 0$, which happens rarely in the one-asset sessions. This suggests that having both assets available to trade has allowed for some degree of incomplete arbitrage against the speculation. However, we still observe prices significantly above 50 (Wilcoxon signed-rank: $p < 0.0001$) in periods where $\delta = 0$, a median price of 57 for Asset A and 59 for Asset B. Furthermore, these above-value (Wilcoxon signed-rank: $p = 0.0002$) median prices are also observed in the initial periods when no information has been revealed in both markets, 56 for Asset A and 57 for Asset B.

Observing prices in complete markets provides an opportunity for an especially simple test of the overpricing hypothesis, by comparing the *sum* of the two assets' prices to 100 in any trading period. Proposition 2 implies that the sum should be greater than the no-arbitrage price of 100 if $\alpha < 5$ and $\beta < 5$. The alternative hypothesis, based on arbitrage pricing, is that the sum of the two prices should not be significantly different from 100.

It is evident from Table IX that, in nearly all cases (19 out of 23), the two asset prices sum to greater than 100. In the $q = \frac{5}{9}$ market, it occurs in all 10 cases. The effect is somewhat muted in the $q = \frac{6}{9}$ treatment, where we observe

TABLE IX
SUM OF MEDIAN PRICES IN COMPLETE
MARKETS SESSIONS^a

δ	$q = \frac{5}{9} (N)$	$q = \frac{6}{9} (N)$
-6		100 (16)
-5		79 (27)
-4		99 (35)
-3	136 (8)	115 (36)
-2	132 (26)	120 (35)
-1	128 (34)	91 (48)
0 (initial period)	129 (73)	104 (65)
0	132 (46)	106 (38)
1	132 (77)	108 (24)
2	124 (72)	108 (15)
3	129 (43)	115 (6)
4	130 (11)	110 (6)
5	132 (5)	
6		101 (2)

^a N = number of transactions.

prices in excess of the no-arbitrage price in 9 of 13 cases. Of possible interest is the observation that all of the exceptions arise when $\delta < 0$ and $q = \frac{6}{9}$. Also worth noting is the fact that the sum of the prices sometimes exceeds 100 by a large amount. In fact, the sum of the prices is 15% or more above the no-arbitrage prices more than half of the time (13 out of 23 cases). The sum of prices is significantly greater than 100 in each treatment and pooled across both treatments ($p < 0.0001$ for $q = \frac{5}{9}$ treatments and both treatments pooled; $p = 0.0226$ for $q = \frac{6}{9}$).

RESULT 7—Complete Markets: Prices in the two-asset markets are systematically higher than no-arbitrage prices. That is, the sum of the prices across the two markets is greater than 100 for nearly all values of δ . This is observed for both treatments.

5.2.2. *Relaxing the Short-Sale Constraint*

Table X presents the aggregate prices for each value of δ for the three sessions in which the short-sale constraint was relaxed, as well as the predicted prices for the homogeneous Bayesian updating model. For 5 out of the 11 values of δ , the median price is actually below the Bayesian price.

Figure 4 shows the disparity between the prices in the sessions with and without the option of buying and selling asset bundles. The median prices in the sessions with a relaxed short-sale constraint are significantly lower than the baseline markets (Mann–Whitney: $p = 0.0037$). In fact, the median price is

TABLE X
 MEDIAN PRICES IN MARKETS WITH SHORT SALES^a

δ	Median Price (N)	Bayesian Price ($\theta = 1$)
-6	14.25 (42)	20.8
-5	22.75 (60)	24.7
-4	34 (56)	29.1
-3	44.5 (121)	33.9
-2	41.75 (181)	39
-1	49 (235)	44.4
0 (initial period)	51.5 (213)	50
0	52 (302)	50
1	55 (310)	55.6
2	65 (191)	61
3	64 (82)	66.1
4	70 (70)	70.9

^a N = number of transactions.

lower for every value of δ except for $\delta = 3$. Allowing short sales substantially reduces speculative overpricing.²⁰

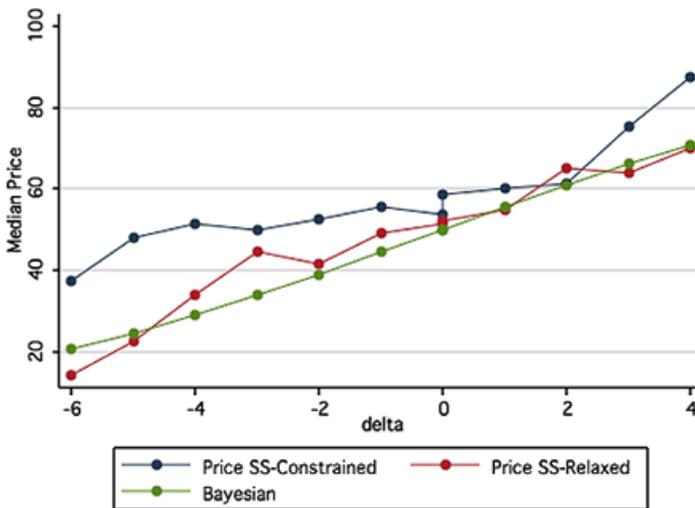


FIGURE 4.—Median prices in short-sales sessions versus Bayesian predictions and markets with no short sales.

²⁰Even though the median prices in these sessions are very close to the Bayesian prices, statistically they are still slightly higher (Wilcoxon signed-rank: $p < 0.0001$).

RESULT 8—Short Sales: Allowing short sales significantly reduces the level of overpricing. Prices are only slightly higher than Bayesian prices.

5.3. Asset Allocations

In addition to properties of equilibrium asset prices, our model also suggests some hypotheses about the dynamics of asset ownership among the traders as a function of the information revealed. We use individual trading data to classify subjects into types based on how their individual holdings vary with δ . One implication of our model for ownership dynamics is that different traders hold the assets over time depending on the pieces of information revealed up to that point. Specifically, when more signals of good news than bad news have been revealed ($\delta > 0$), the θ_{\max} traders are the optimists and should be net buyers of the asset, while everyone else should sell the asset. On the other hand, when more signals of bad news than good news have been revealed ($\delta < 0$), the θ_{\min} traders are the optimists and should hold the asset. To investigate these predicted switches, we compare the distribution of asset holdings across traders in periods where $\delta > 0$ to the holdings distribution in periods where $\delta < 0$.

We categorize each trader into one of behavioral types based on whether his/her net holdings (end-of-period holdings minus initial endowment) is positive or negative. For each trader, we do this separately for $\delta > 0$ periods and $\delta < 0$ periods.²¹ For each of these two ranges of δ , a trader is counted as having zero net holdings if a trader's mean net holdings are less than the standard error of that trader's net holdings over that range of δ . Otherwise the trader is counted as having either positive or negative net holdings over that range of δ . Traders with positive net holdings when $\delta > 0$ and negative or zero net holdings when $\delta < 0$, or zero net holdings when $\delta > 0$ and negative net holdings when $\delta < 0$, are categorized as *Gullible* or θ_{\max} types because they are net buyers only when there is more good news than bad news. Traders with positive holdings when $\delta < 0$ and negative or zero net holdings when $\delta > 0$, or zero net holdings when $\delta < 0$ and negative net holdings when $\delta > 0$, are categorized as *Skeptical* or θ_{\min} types because they are net buyers only when there is more bad news than good news. *Always Sell* types have negative or zero net holdings in both $\delta > 0$ and $\delta < 0$ periods, and these correspond in the model to traders with intermediate values of θ , with $\theta_{\min} < \theta < \theta_{\max}$.

Table XI summarizes the distribution of trader types across treatments. The vast majority of traders (82.4%) in our markets are categorized as *Gullible*, *Skeptical*, or *Always Sell* types, which is consistent with the heterogeneous beliefs model. There is a small residual category of traders who do not fall in either of these three categories, and we refer to them as *Noise Traders*. These few traders, who have nonnegative and sometimes positive net holdings for both ranges of δ , are difficult to reconcile with the existing model.

²¹Periods where $\delta = 0$ are not included in the analysis because the model makes no prediction about asset holdings in these periods.

TABLE XI
DISTRIBUTION OF TRADER TYPES

Treatment	Always Sell	Gullible	Skeptical	Always Buy	Overall
5/9	10	9	10	5	34
6/9	10	8	9	7	34
Total	20	17	19	12	68

RESULT 9—Trader Types: Most traders are classified in one of the three categories: fickle, skeptical, and always sell, corresponding to high, low, and intermediate values of θ , respectively.

6. CONCLUSION

We study pricing in asset markets with public information flows and the short-sale constraint when traders have heterogeneous beliefs. We analyze a simple parsimonious model of such heterogeneity with a single parameter that indexes whether a trader overweights or underweights new information relative to Bayesian updating. Building on [Harrison and Kreps \(1978\)](#) and [Morris \(1996\)](#), this model generates equilibrium price dynamics that exhibit two key properties: speculative overpricing and asymmetric response to good and bad news.

We report data from a series of laboratory markets for an asset whose terminal payoff is contingent upon an unknown state of the world. At regular intervals during the trading, a sequence of ten informative but imperfect signals are publicly revealed to the traders at regular intervals. We find asset prices consistently above what would be predicted by a model of homogeneous Bayesian updating to a common posterior. Theoretically, an important component of these high prices is speculative overpricing, which takes the form of a speculative premium equal to the difference between the equilibrium price and the maximum hold-to-maturity valuation among all traders. Traders are willing to pay more for the asset than its “true” valuation as long as there is a positive probability that some other trader may value it even more after some future sequence of signals.

We also document an asymmetric price response to information in good news versus bad news regimes, in that prices respond more strongly to information when good news has outweighed bad news. This arises in the model as a direct consequence of belief heterogeneity, because low- θ type traders who update too conservatively have the most optimistic valuation in bad news regimes; this valuation then determines the price, thus dampening the effect of information. The high- θ types, who over-react to public information, have the highest valuation when there has been more good news than bad news; thus prices reflect this over-reaction.

To measure the speculative component of overpricing and to test explicitly for the presence of belief heterogeneity, we estimate the maximum and minimum value of θ in each session, and find significant heterogeneity of beliefs in all but one of the baseline sessions. The estimated distribution of θ types implies a quantitative measure of the speculative premium in each period of each market. We find that the estimated speculative premium is significantly positive in periods with no permanent optimist, as predicted. In contrast, the estimated speculative premium is significantly lower when there is a permanent optimist, and is usually not significantly different from zero.

Next, we ask whether the speculative overpricing can be curbed institutionally. To answer this question, we ran additional sessions where we manipulated the market organization in two different ways: (a) by instituting complete markets with the full set of two Arrow–Debreu securities being traded in parallel, one market for asset A and one for asset B; and (b) by instituting a market where short sales are allowed. In the complete markets environment, significant overpricing persists: the transaction prices of the two assets add up to more than the certain payoff of holding one unit of each asset, 100. Just as in the one-market sessions, we observe prices that are significantly higher than 50 in the initial periods when no information has been revealed, as well as in periods where the number of good news signals equals the number of bad news signals.

In contrast, when the short-sale constraint is removed, prices are significantly reduced and are close to the Bayesian benchmark. Overall, our results demonstrate that the short-sale constraint can be an important factor leading to speculative overpricing in asset markets, which may have implications for policies that explicitly limit the extent to which traders may engage in short selling.

Finally, the model also generates predictions about patterns of asset ownership that depend on the sequence of public signals. To study the predictions about individual asset ownership, we categorize traders into several categories depending on whether their trading behavior is consistent with being a high θ -type (net buyer with good news, net seller with bad news), a low- θ type (net buyer with bad news, net seller with good news), or an intermediate- θ type (net seller). Most traders' net holdings patterns are consistent with the heterogeneous- θ model, with two caveats: most traders do not reduce their holdings to exactly zero when they are net sellers; and a few traders average positive net holdings in both good news regimes and bad news regimes.

We conclude that the heterogeneous beliefs model of asset pricing is broadly supported by our data. Methodologically, the “public information flow” asset market design used here is an innovation to laboratory markets that makes it possible to address important theoretical questions about asset pricing dynamics. There are a variety of different directions to take this work and our findings are suggestive of some interesting theoretical and experimental extensions. On the theoretical side, one could enrich the type space by considering multidimensional time-dependent types where θ_{it} varies over time. One

could also consider alternative specifications of belief heterogeneity. It would be useful to extend both the Morris model and our model to include risk-averse traders or to incorporate private information. In principle, one would expect the qualitative properties of the dynamic trajectory of asset prices (speculative premium, asymmetry, and horizon effects) to continue to hold in these more general models, but the holdings predictions would not be as extreme. However, until these difficult theoretical problems are resolved, we leave this as a conjecture.

APPENDIX: PROOFS

PROOF OF PROPOSITION 1: The posterior belief that $\varpi = A$ for traders of type θ , given a sequence of good and bad news announcements $(\alpha, \beta = t - \alpha)$, is

$$\rho_{it}^\alpha(\theta) = \frac{q^{\alpha\theta}(1 - q)^{\beta\theta}}{q^{\alpha\theta}(1 - q)^{\beta\theta} + q^{\beta\theta}(1 - q)^{\alpha\theta}}.$$

Since $q > 0.5$,

$$\theta_{\max} = \arg \max_{i \in \mathcal{I}} \{\rho_{it}^\alpha(\theta_i)\} \quad \text{if } \delta > 0,$$

and

$$\theta_{\min} = \arg \max_{i \in \mathcal{I}} \{\rho_{it}^\alpha(\theta_i)\} \quad \text{if } \delta < 0.$$

Traders with the highest θ place the greatest weight on signals, so their posterior for state A will be highest of all traders when $\alpha > \beta$. On the other hand, traders with the lowest θ underweight the signals the most, so their posterior for state A will be highest of the traders when $\beta > \alpha$. When $\alpha - \beta = \delta = 0$, then $\rho_{it}(\alpha) = p = 0.5 \forall i$, that is, all traders' beliefs coincide, and every trader is a current optimist.

To prove (i), consider period t and any sequence such that $|\delta| + t < T$. Thus, if all future signals are a -signals (i.e., $s_{t+1} = \dots = s_T = a$), then the current optimist at period T is a θ_{\max} trader. Similarly, if all future signals are b -signals (i.e., $s_{t+1} = \dots = s_T = b$), then the current optimist at period T is a θ_{\min} trader. Therefore, there is no permanent optimist at period t . That is, there is no permanent optimist at t if and only if it is uncertain which trader type will be the current optimist in the final period, T .

To complete the proof, we need to show that the speculative premium is positive. Suppose, without loss of generality, that in period t there has been a sequence of signals $\{s_1, \dots, s_t\}$ with $0 \leq \delta < T - t$, so a θ_{\max} trader is a current optimist. (The argument is the same for the case of $0 \leq -\delta < T - t$ and θ_{\min} is a current optimist.) Let $t' = t + 1 + \alpha - \beta$. Note that $t' \leq T$ since $\alpha - \beta + t < T$.

Consider all sequences of signals $s_t' = \{s_{t+1}, \dots, s_{t'}\}$. For exactly one such sequence, $s_t' = (b, \dots, b)$, the sequence with all b signals, a θ_{\max} trader is *not* a current optimist at t' . (Instead, a θ_{\min} trader will be the current optimist.) The equilibrium price following this sequence is just $P_{t'}(\alpha)$. Because θ_{\max} is *not* a current optimist at t' for this sequence, it implies that $P_{t'}(\alpha) > \rho_{t'}^\alpha(\theta_{\max})$. For all other sequences $s_t' \neq (b, \dots, b)$, a θ_{\max} trader is still a current optimist at t' . Let $z(s_t')$ denote the θ_{\max} trader's belief at t that the exact sequence of signals from period $t + 1$ to period t' will be s_t' , and denote by $\tilde{\alpha}(s_t')$ the total number of a signals out of *all* signals $s_1, \dots, s_{t'}$, given that α of the first t signals were a signals. Thus, the θ_{\max} trader's belief at t that the state of the world is A equals $\rho_{t'}^{\tilde{\alpha}(s_t')}(\theta_{\max})$. Because θ_{\max} is a current optimist at t' , it implies that $P_{t'}(\tilde{\alpha}(s_t')) \geq \rho_{t'}^{\tilde{\alpha}(s_t')}(\theta_{\max})$ for all $s_t' \neq (b, \dots, b)$. However, by the recursive definition of prices, and because θ_{\max} is a current optimist for *every* period from t to $t' - 1$ (since we constructed $t' = t + 1 + \alpha - \beta$), the current price is given by

$$\begin{aligned} P_t(\alpha) &= z(b, \dots, b)P_{t'}(\alpha) + \sum_{s_t' \neq (b, \dots, b)} z(s_t')P_{t'}(\tilde{\alpha}(s_t')) \\ &\geq z(b, \dots, b)P_{t'}(\alpha) + \sum_{s_t' \neq (b, \dots, b)} z(s_t')\rho_{t'}^{\tilde{\alpha}(s_t')}(\theta_{\max}) \\ &> z(b, \dots, b)\rho_{t'}^\alpha(\theta_{\max}) + \sum_{s_t' \neq (b, \dots, b)} z(s_t')\rho_{t'}^{\tilde{\alpha}(s_t')}(\theta_{\max}) \\ &= \rho_t^\alpha(\theta_{\max}). \end{aligned}$$

The last line follows from the fact that, given the updating formulas, $\rho_t^\alpha(\theta_{\max}) = \sum_{s_t'} z(s_t')\rho_{t'}^{\tilde{\alpha}(s_t')}(\theta_{\max})$. Hence $P_t(\alpha) > \rho_t^\alpha(\theta_{\max})$, and there is a positive speculative premium when there is no permanent optimist.

We now move to the proof of (ii). If, at some t , we have $\alpha - \beta \geq T - t$, then $\alpha \geq \beta$ and hence $\rho_t^\alpha(\theta_{\max}) \geq \rho_{it}^\alpha(\theta_i) \forall i$. Furthermore, for all $t' = t + 1, \dots, T$ and all possible sequences of signals $\{s_{t+1}, \dots, s_{t'}\}$, $\alpha \geq t' - \frac{T}{2} \geq t' - \alpha = \beta$, θ_{\max} will be the current optimist: $\rho_{t'}^{\theta_{\max}}(\alpha) \geq \rho_{it'}(\alpha) \forall i$. Thus a θ_{\max} trader is a permanent optimist at t following any sequence of signals $\{s_1, \dots, s_t\}$ such that $|s \in \{s_1, \dots, s_t\} | s = a | = \alpha \geq \frac{T}{2}$. By a similar argument, a θ_{\min} trader is a permanent optimist at t following any sequence of signals $\{s_1, \dots, s_t\}$ such that $|s \in \{s_1, \dots, s_t\} | s = b | \geq \frac{T}{2}$.

To complete the proof of (ii), let τ index the number of periods left until the end of the market. We prove by induction on τ that if θ_{\max} is a permanent optimist at t , then, for all possible continuation sequences of signals up to period $T - \tau$ (i.e., for all $\alpha' = \alpha, \alpha + 1, \dots, \alpha + T - t - \tau$), $P_\tau(\alpha') = \rho_\tau^{\alpha'}(\theta_{\max})$, $\tau = 0, \dots, T - t$. First, note that this is trivially true for $\tau = 0$, since this means

$t = T$, and so it is the last period. The last period price is given by $P_T(\alpha') = \bar{p}_T(\alpha') = \rho_T^{\alpha'}(\theta_{\max})$. The last holds because θ_{\max} is a permanent optimist at t and $\alpha' \geq \alpha \geq \frac{t}{2}$. Next we show that, if $0 < \tau < T - t$ and $P_{T-\tau}(\alpha') = \rho_{T-\tau}^{\alpha'}(\theta_{\max})$ for all $\alpha' = \alpha, \alpha + 1, \dots, \alpha + T - t - \tau$, then $P_{T-(\tau+1)}(\alpha') = \rho_{T-(\tau+1)}^{\alpha'}(\theta_{\max})$ for all $\alpha' = \alpha, \alpha + 1, \dots, \alpha + T - t - (\tau + 1)$. Fix some $\alpha' \in \{\alpha, \alpha + 1, \dots, \alpha + T - t - (\tau + 1)\}$. By definition, in period $T - (\tau + 1)$, the equilibrium price is

$$\begin{aligned} P_{T-(\tau+1)}(\alpha') &= \varphi_{T-(\tau+1)}(\alpha')P_{T-\tau}(\alpha' + 1) + (1 - \varphi_{T-(\tau+1)}(\alpha'))P_{T-\tau}(\alpha') \\ &= \varphi_{T-(\tau+1)}(\alpha')\rho_{T-\tau}^{\alpha'+1}(\theta_{\max}) + (1 - \varphi_{T-(\tau+1)}(\alpha'))\rho_{T-\tau}^{\alpha'}(\theta_{\max}) \\ &= \rho_{T-(\tau+1)}^{\alpha'}(\theta_{\max}). \end{aligned} \quad Q.E.D.$$

PROOF OF PROPOSITION 2: As before, let τ index the number of periods left until the end of the market. We again prove the result by induction on τ . First, we show the proposition is true for $\tau = 0$ and $\tau = 1$. Note that since $\alpha > \frac{t}{2}$, $\rho_t^*(\alpha) = \rho_t^{\theta_{\max}}(\alpha)$ because $\delta > 0$. Similarly, $\rho_t^*(t - \alpha) = \rho_t^{\theta_{\min}}(t - \alpha) > \rho_t^{\theta_{\max}}(t - \alpha)$. Since $\rho_t^{\theta_{\max}}(\alpha) + \rho_t^{\theta_{\max}}(t - \alpha) = 1$, $\rho_t^{\theta_{\min}}(t - \alpha) > 1 - \rho_t^{\theta_{\max}}(\alpha)$. If $\tau = 0$ or 1 , $t \geq T - 1$. Thus, $\alpha > \frac{t}{2}$ implies $\alpha \geq \frac{T}{2}$. By Proposition 1, if $\alpha \geq \frac{T}{2}$, then $P_t(\alpha) = \rho_t^*(\alpha)$ and $P_t(t - \alpha) = \rho_t^*(t - \alpha)$. Therefore

$$\begin{aligned} P_t(t - \alpha) &= \rho_t^{\theta_{\min}}(t - \alpha) \\ &> (1 - \rho_t^{\theta_{\max}}(\alpha)) = (1 - P_t(\alpha)). \end{aligned}$$

Note that, given the definition of τ , if τ corresponds to $T - t$, then $\tau + 1$ corresponds to $T - t - 1$. To complete the proof, letting $t = T - \tau$, we show that, for any $T - 2 > \tau > 1$ and for all $t > \alpha > \frac{t}{2}$, $1 - P_t(\alpha) < P_t(t - \alpha)$ implies $1 - P_{t-1}(\alpha) < P_{t-1}(t - 1 - \alpha)$. By the equilibrium pricing equation,

$$\begin{aligned} P_{t-1}(\alpha) &= \varphi_{t-1}(\alpha)P_t(\alpha + 1) + (1 - \varphi_{t-1}(\alpha))P_t(\alpha) \\ &= \varphi_{t-1}^{\theta_{\max}}(\alpha)P_t(\alpha + 1) + (1 - \varphi_{t-1}^{\theta_{\max}}(\alpha))P_t(\alpha). \end{aligned}$$

The second line follows because $\delta > 0$. The equilibrium pricing equation after a sequence with equal number of signals ($t - 1$) but the numbers for the pieces of good news and bad news reversed is

$$\begin{aligned} P_{t-1}(t - 1 - \alpha) &= \varphi_{t-1}(t - 1 - \alpha)P_t(t - \alpha) \\ &\quad + (1 - \varphi_{t-1}(t - 1 - \alpha))P_t(t - 1 - \alpha) \\ &= \varphi_{t-1}^{\theta_{\min}}(t - 1 - \alpha)P_t(t - \alpha) \\ &\quad + (1 - \varphi_{t-1}^{\theta_{\min}}(t - 1 - \alpha))P_t(t - 1 - \alpha). \end{aligned}$$

The second line is due to the fact that $t - 1 - \alpha < \frac{t}{2}$, so the number of a signals is less than the number of b signals.

Because the proposition is assumed to be true for $\tau = T - t$, we have $1 - P_t(\alpha + 1) < P_t(t - (\alpha + 1)) = P_t(t - 1 - \alpha)$ and $1 - P_t(\alpha) < P_t(t - \alpha)$. Thus

$$\begin{aligned} 1 - P_{t-1}(\alpha) &= \varphi_{t-1}^{\theta_{\max}}(\alpha)[1 - P_t(\alpha + 1)] \\ &\quad + (1 - \varphi_{t-1}^{\theta_{\max}}(\alpha))[1 - P_t(\alpha)] \\ &< \varphi_{t-1}^{\theta_{\max}}(\alpha)P_\tau(t - 1 - \alpha) + (1 - \varphi_{t-1}^{\theta_{\max}}(\alpha))P_t(t - \alpha) \\ &< (1 - \varphi_{t-1}^{\theta_{\min}}(t - 1 - \alpha))P_\tau(t - 1 - \alpha) \\ &\quad + \varphi_{t-1}^{\theta_{\min}}(t - 1 - \alpha)P_t(t - \alpha) \\ &= P_{t-1}(t - 1 - \alpha), \end{aligned}$$

because $P_\tau(t - 1 - \alpha) < P_t(t - \alpha)$ and $\varphi_{t-1}^{\theta_{\max}}(\alpha) > 1 - \varphi_{t-1}^{\theta_{\min}}(t - 1 - \alpha)$, the latter following from $\theta_{\max} > \theta_{\min}$. *Q.E.D.*

PROOF OF PROPOSITION 3: Again τ indexes the number of periods left until the end of the market. We prove by induction that $\pi_{T-\tau}(\alpha) - \pi_{T-\tau-2}(\alpha - 1) \leq 0$ for all $\tau = 0, \dots, T - 2$ and for all $\alpha = 1, \dots, T - \tau - 1$.

We first show that this is true for $\tau = 0$. By Proposition 1, $\pi_{\tau=0}(\alpha) = \pi_T(\alpha) = 0$ and $\pi_{T-\tau-2}(\alpha - 1) = \pi_{T-2}(\alpha - 1) \geq 0$. Thus $\pi_{T-\tau}(\alpha) - \pi_{T-\tau-2}(\alpha - 1) \geq 0$ when $\tau = 0$.

Note that, given the definition of τ , if τ corresponds to t , then $\tau + 1$ corresponds to $t - 1$. To complete the proof, letting $t = T - \tau > 2$, we need to show that if $\pi_t(\alpha) \leq \pi_{t-2}(\alpha - 1)$ for all $\alpha = 1, \dots, t - 1$, then $\pi_{t-1}(\alpha) \leq \pi_{t-3}(\alpha - 1)$ for all $\alpha = 1, \dots, t - 2$. To prove this, pick any $\alpha \in \{1, \dots, t - 2\}$. We have

$$\begin{aligned} \pi_{t-1}(\alpha) &= \varphi_{t-1}^*(\alpha)[\pi_t(\alpha + 1) + \rho_t^*(\alpha + 1)] \\ &\quad + (1 - \varphi_{t-1}^*(\alpha))[\pi_t(\alpha) + \rho_t^*(\alpha)] - \rho_{t-1}^*(\alpha) \\ &= \pi_t(\alpha) + \rho_t^*(\alpha) - \rho_{t-1}^*(\alpha) \\ &\quad + \varphi_{t-1}^*(\alpha)[\pi_t(\alpha + 1) - \pi_t(\alpha) + \rho_t^*(\alpha + 1) - \rho_t^*(\alpha)], \end{aligned}$$

while the speculative premium for a sequence with the same δ but two more periods left before the end of the market, $\pi_{t-3}(\alpha - 1)$, is

$$\begin{aligned} \pi_{t-3}(\alpha - 1) &= \varphi_{t-3}^*(\alpha - 1)[\pi_{t-2}(\alpha) + \rho_{t-2}^*(\alpha)] - \rho_{t-3}^*(\alpha - 1) \\ &\quad + (1 - \varphi_{t-3}^*(\alpha - 1))[\pi_{t-2}(\alpha - 1) + \rho_{t-2}^*(\alpha - 1)] \\ &= \pi_{t-2}(\alpha - 1) + \rho_{t-2}^*(\alpha - 1) - \rho_{t-3}^*(\alpha - 1) \\ &\quad + \varphi_{t-3}^*(\alpha - 1) \\ &\quad \times [\pi_{t-2}(\alpha) - \pi_{t-2}(\alpha - 1) + \rho_{t-2}^*(\alpha) - \rho_{t-2}^*(\alpha - 1)]. \end{aligned}$$

We now calculate the difference between these two speculative premiums, $\pi_{t-3}(\alpha - 1) - \pi_{t-1}(\alpha)$. A number of terms cancel out in this difference because $\rho_t^*(\alpha) = \rho_{t-2}^*(\alpha - 1)$, $\rho_{t-1}^*(\alpha) = \rho_{t-3}^*(\alpha - 1)$, $\rho_t^*(\alpha + 1) = \rho_{t-2}^*(\alpha)$, and $\varphi_{t-1}^*(\alpha) = \varphi_{t-3}^*(\alpha - 1) = \tilde{\varphi}^* \in (0, 1)$. Thus we are left with

$$\begin{aligned} \pi_{t-3}(\alpha - 1) - \pi_{t-1}(\alpha) &= (1 - \tilde{\varphi}^*)(\pi_{t-2}(\alpha - 1) - \pi_t(\alpha)) \\ &\quad + \tilde{\varphi}^*(\pi_{t-2}(\alpha) - \pi_t(\alpha + 1)). \end{aligned}$$

We know that $\pi_{t-2}(\alpha - 1) - \pi_t(\alpha) \geq 0$ and $\pi_{t-2}(\alpha) - \pi_t(\alpha + 1) \geq 0$ by the induction hypothesis. Therefore, since $\tilde{\varphi}^* \in (0, 1)$, we get $\pi_{t-3}(\alpha - 1) - \pi_{t-1}(\alpha - 1) \geq 0$. Q.E.D.

REFERENCES

- ANDERSON, M., AND S. SUNDER (1995): "Professional Traders as Intuitive Bayesians," *Organizational Behavior and Human Decision Processes*, 64, 185–202. [1945]
- ASPAROUHOVA, E., P. BOSSAERTS, J. EGUIA, AND W. R. ZAME (2009): "Asset Pricing and Asymmetric Reasoning," Working Paper, University of Utah, Caltech, NYU and UCLA. [1941,1943]
- BIAIS, B., AND P. BOSSAERTS (1998): "Asset Prices and Trading Volume in a Beauty Contest," *Review of Economic Studies*, 65, 307–340. [1942]
- BOSSAERTS, P., C. R. PLOTT, AND W. R. ZAME (2007): "Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments," *Econometrica*, 75, 993–1038. [1941]
- CAMERER, C. (2003): *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton, NJ: Princeton University Press. [1945]
- EL-GAMAL, M. A., AND D. M. GREYER (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1137–1145. [1938,1945]
- FORSYTHE, R., T. R. PALFREY, AND C. R. PLOTT (1982): "Asset Valuation in an Experimental Market," *Econometrica*, 50, 537–568. [1940]
- (1984): "Futures Markets and Information Efficiency: A Laboratory Examination," *Journal of Finance*, 39, 955–981. [1940]
- FRIEDMAN, D., G. W. HARRISON, AND J. W. SALMON (1984): "The Informational Efficiency of Experimental Asset Markets," *Journal of Political Economy*, 92, 349–408. [1940]
- GOEREE, J., T. R. PALFREY, B. W. ROGERS, AND R. MCKELVEY (2007): "Self-Correcting Information Cascades," *Review of Economic Studies*, 74, 733–762. [1945]
- GROSSMAN, S. J., AND J. E. STIGLITZ (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70 (3), 393–408. [1941]
- HARRIS, M., AND A. RAVIV (1993): "Differences of Opinion Make a Horse Race," *Review of Financial Studies*, 6 (3), 473–506. [1942]
- HARRISON, J. M., AND D. M. KREPS (1978): "Speculative Investor Behavior in a Stock Market With Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323–336. [1937,1969]
- MILLER, E. (1977): "Risk, Uncertainty and Divergence of Opinion," *Journal of Finance*, 32, 1151–1168. [1942]
- MORRIS, S. (1996): "Speculative Investor Behavior and Learning," *Quarterly Journal of Economics*, 111, 1111–1133. [1942,1949,1969]
- OTTAVIANI, M., AND P. N. SORENSEN (2007): "Aggregation of Information and Beliefs in Prediction Markets," Working Paper, Northwestern University and University of Copenhagen. [1942,1943]
- PALFREY, T. R., AND S. W. WANG (2012): "Supplement to 'Speculative Overpricing in Asset Markets With Information Flows'," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/8781_Instructions_to_experimental_subjects.pdf; http://www.econometricsociety.org/ecta/Supmat/8781_data_and_programs.zip. [1951]

- PLOTT, C. R., AND S. SUNDER (1982): "Efficiency of Experimental Security Markets With Insider Information: An Application of Rational Expectations Models," *Journal of Political Economy*, 90, 663–698. [1941]
- (1988): "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56, 1085–1118. [1939,1941]
- PORTER, D. P., AND V. L. SMITH (2003): "Stock Market Bubbles in the Laboratory," *Journal of Behavioral Finance*, 4, 7–20. [1942]
- RADNER, R. (1979): "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, 47 (3), 655–678. [1941]
- SCHEINKMAN, J. A., AND W. XIONG (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111, 1183–1220. [1942]
- SMITH, V. L., G. SUCHANEK, AND A. WILLIAMS (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56 (5), 1119–1151. [1941]
- SUNDER, S. (1995): "Experimental Asset Markets: A Survey," in *Handbook of Experimental Economics*, ed. by A. Roth and J. Kagel. Princeton, NJ: Princeton University Press. [1941]

*Division of the Humanities and Social Sciences, Caltech, Pasadena, CA 91125,
U.S.A.; trp@hss.caltech.edu*

and

*Dept. of Economics, University of Pittsburgh, Pittsburgh, PA 15260, U.S.A.;
swwang@pitt.edu.*

Manuscript received August, 2009; final revision received March, 2012.