Rational and Heuristic-Driven Trading Panics in an Experimental Asset Market

Chad Kendall*

April 7, 2015

Abstract

In financial markets, potential adverse price movements may induce traders to rush to trade as soon as possible. On the other hand, time-consuming research produces better asset value information. I study this trade-off both theoretically and experimentally. I derive conditions under which equilibrium forces cause traders to rationally panic, trading simultaneously as each attempts to front-run the others. In the laboratory data, robust, rational panics occur, resulting in poor information aggregation as traders forgo better information. Panics that cannot be explained by equilibrium behavior are also frequent, further hampering information aggregation, and resulting in trade clustering and positive short-term correlation in returns. These real-world phenomena are a result of about 40% of traders relying on a simple trading heuristic that experience does not eliminate. I discuss evidence that the heuristic may be driven by prospect theory preferences, and its implications for asset markets outside of the laboratory.

1 Introduction

In financial markets, traders must decide how much information to acquire before trading an asset. Obtaining better information leads to higher expected profits and more informative

*Department of Finance and Business Economics, Marshall School of Business, University of Southern California, 3670 Trousdale Parkway, Ste. 308, BRI-308, MC-0804, Los Angeles, CA 90089-0804 (e-mail chadkend@marshall.usc.edu). I would like to thank Andrea Frazzini, Francesco Trebbi, Patrick Francois, and Ryan Oprea, for their substantial guidance and feedback. I also benefited from comments from Andreas Park, Sophie Moinas, Yoram Halevy, Vadim Marmer, as well as participants in various seminars: UBC and USC brown bag lunches, Miami Behavioral Finance Conference, and Cubist Systematic Strategies. I gratefully acknowledge support from SSHRC through a Joseph-Armand Bombardier CGS scholarship.

asset prices. However, when information can only be acquired over time, it comes at a cost: if others trade, expected prices move adversely, reducing one's trading profit.[1] This fear of adverse price movements can cause a rational panic - defined as a trading episode in which traders simultaneously rush to trade before receiving full information. In a rational panic, asset values poorly aggregate information as each trader forgoes the opportunity to acquire better information. This paper develops and tests, in a laboratory setting, a model of panics and their associated informational consequences.[2]

Why study market panics in a laboratory environment? Real-world evidence is difficult to interpret because we do not typically observe the information traders acquire nor when they acquire it.[3] By overcoming these issues, a laboratory experiment can provide insight into two main questions of interest. First, are panics consistent with equilibrium behavior or do they result from other rules-of-thumb? Second, how do panics (rational or otherwise) affect the informational content of prices?

I introduce a novel model of trading with two important features. First, each trader in the model has the opportunity to trade in many trading periods, but is restricted to only one trade. Unlike models in which multiple trades are allowed, this restriction allows panics (simultaneous trades) to be easily identified. Second, traders receive private information over time such that they face a trade-off between better information and the desire to trade before others. Specifically, traders receive initial (poor quality) information when they arrive to the market (simultaneously), and are then given several opportunities to trade before receiving additional information just prior to the final trading period.

The common-value asset and signals in the model are binary. Risk-neutral traders trade with a market maker who posts a single price equal to the expected value of the asset conditional on all public information.[4] Given their private information, traders can earn an

[1]Unconditional correlation in signals means prices move against you in expectation: when one has good information about a stock, others are likely to have similar information.

[2]'Panics' defined in this way are distinct from 'crashes' and 'frenzies'. In the same way that a panic in a crowded theater may or may not lead to people being trampled to death, a panic may or may not lead to a price crash. That is, panics are caused by actions taken out of fear (in this case of adverse price movements), and may lead to one of several consequences. A price crash is one such consequence, but informational losses are another, which I study here. For a concrete example of a panic that does not lead to a price crash, consider the fact that earnings announcements generate spikes in volume (Frazzini and Lamont (2006)), but not necessarily large price movements. For papers that study (theoretically) how rational panics can lead to crashes, see Romer (1993), Bulow and Klemperer (1994), Smith (1997), Lee (1998), Barlevy and Veronesi (2003), Brunnermeier and Pedersen (2005), and Pedersen (2009).

[3]Methods of inferring information from trades do exist. For example, Hasbrouck (1991) uses the persistence of price impacts and Easley et al. (1996) construct a structural model based on the arrival of news. Kendall (2015) discusses the theoretical implications of traders rushing to trade for such indirect evidence.

[4]The model is in the spirit of Glosten and Milgrom (1985), but I abstract from the adverse selection problem between the market maker and the informed traders. Here, the market maker provides liquidity so that noise traders are not necessary. As shown in Kendall (2015), incorporating adverse selection by having

expected profit by trading. I show that the standard result in which traders buy with favorable private information and sell with unfavorable private information applies. In the main theoretical contribution, I provide precise predictions about when trades occur in equilibrium.[5] I provide a pair of sufficient conditions, one of which ensures that the equilibrium prediction is a rational panic (all traders trade in the first period), and the other of which ensures that the prediction is for traders to wait for better information, avoiding a panic (trade in the final period). As one may expect, the main driver of whether one should panic or wait is the difference between initial and final signal qualities. When this difference is large, waiting intuitively becomes more valuable. However, the signal qualities also affect the information revealed by trades and therefore the cost side of the trade-off (prices move more when others have better information). The theoretical results provide sufficient conditions for the value of information to dominate and for the endogenous cost to dominate.

The equilibrium timing predictions are quite stark: the time at which one should trade does not depend upon the price or one's private information. Any trends in prices are therefore irrelevant. However, a natural intuition is that as prices rise such that the asset is more likely to be good, one may want to buy the asset immediately. Buying an asset that is more likely to be good is compelling, and one may reason that waiting simply results in paying a higher price for it. While partially correct, this intuition fails because an asset that is more likely to be good is already priced to account for this fact. Instead, expected profits only depend upon the difference between public and private beliefs, so that one should either trade according to one's current private information or wait for more of it. By allowing for many trading periods so that price trends may form, I designed the model such that this faulty intuition, if used, is observable in the data.

In the experiment, I conduct two treatments with parameters such that rational panics should occur in the first, but not in the second. In the treatment in which panics are rational,

---

the market maker post separate bid and ask prices makes analysis of the timing decisions more complicated, but doesn't affect the main trade-off between better information and costly price moves which is the main object of interest.

[5]Precise theoretical predictions are mostly absent in experimental studies of markets with endogenous trade timing. Park and Sgroi (2012), in the most closely related paper, provide qualitative predictions guided by theory in an endogenous timing setting. Shachat and Srinivasan (2011) study trading with sequential arrival of information, but in an environment in which a no trade theorem applies. Bloomfield et al. (2005) study the choice between market and limit orders in the absence of theoretical predictions. Several papers consider timing decisions in environments in which prices are fixed and theoretical predictions are known. See, for example, Sgroi (2003) and Ziegelmeyer et al. (2005) who each implement the irreversible investment model of Chamley and Gale (1994). Ivanov et al. (2009,2013) and Çelen and Hyndham (2012) study similar environments. Incentives to wait are quite different in these papers as observing others' decisions can be beneficial when their information is not incorporated into prices. Finally, Brunnermeier and Morgan (2010) study timing decisions theoretically and experimentally in a game related to both preemption and war of attrition games.

they always occur, providing the first available laboratory evidence that subjects rationally rush to trade to avoid adverse price movements and suggesting that theory can be a useful guide to predicting panics in the field. Subjects forgo acquiring perfect information about the asset value 99.7% of the time, instead rationally choosing to trade immediately on lower quality information. 75% of trades occur in the period of arrival, leading to extreme clustering of trades as subjects rationally rush to be the first to trade. This treatment provides a robust demonstration of rational panics and the corresponding informational consequences established in the theory.

In the treatment in which subjects should rationally wait for additional information, the intuitive behavior described previously is much more common than rational behavior. To analyze this behavior, I formally define a heuristic motivated by the intuition and confirm its predictions. The heuristic produces non-equilibrium panics (trade clustering) that result in informational losses not predicted by theory. It also produces short-term positive correlations in returns due to the fact that traders chase price trends. Chasing price trends can lead to informational herding (trading against one's private information), but perhaps surprisingly, the positive correlation in returns is *not* driven primarily by herding.[6] Instead, the correlation is driven by traders with different private information choosing to trade at different times. At the individual level, about one third of subjects clearly follow the heuristic, a much higher percentage than behave rationally. Given the large proportion of subjects following the heuristic and the emergence of positive correlations in returns and trade clustering, two well known features of real markets, it seems plausible that the heuristic plays a role in the field.[7]

To study the heuristic predictions in greater detail and determine its robustness, I develop a second model that I refer to as the *Diagnostic* model. One difficulty with interpreting the data of the main model is that selection issues arise: subjects that trade early on exit the sample so that one cannot observe what they would have done at other price levels. To overcome this issue, in the Diagnostic model, traders arrive sequentially and make simple,

[6]The literature on informational herding began with the papers of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). Various definitions of herding and its counterpart, contrarianism, are present in the literature. I follow the definitions of Avery and Zemsky (1998) and the subsequent experimental literature that tests their model (Cipriani and Guarino (2005) and Drehmann et al. (2005)). See Section 3.3 for the formal definitions. For a survey on herding in financial markets, see Devenow and Welch (1996). Herding has has long been considered an explanation for correlation in returns. See Hirshleifer and Teoh (2003) for an extensive review.

[7]Dufour and Engle (2000) provide empirical evidence of trade clustering. Positive short-term correlation in returns is often attributed to underreaction and has received considerable attention both theoretically and empirically. For theoretical explanations, see Barberis et al. (1998), Daniel et al. (1998), and Hong and Stein (1999). For a review of empirical evidence, see Daniel et al. (1998). Kirchler (2009) demonstrates underreaction in an experimental double-auction environment in which asset values follow a stochastic process.

binary timing decisions. The model is a variation of the model of Kendall (2015) in which subjects trade in an overlapping sequence such that an additional trader arrives before one's second trading opportunity.[8]

I construct a pair of treatments for the Diagnostic model, showing theoretically that in one it is rational to rush to trade upon arrival to the market, and in the other it is rational to wait for more information. The prediction of the heuristic, on the other hand, is that subjects that arrive at times when the price is high should be more likely to rush to buy. In this way, the Diagnostic model gives a clean prediction: the higher the price, the more likely a subject should be to rush.

Comparing behavior across the two Diagnostic treatments, more rushing is observed in the treatment in which it is rational to rush. However, frequent departures from rationality are observed in both treatments. The heuristic prediction of more rushing at higher prices is confirmed in both treatments. Furthermore, at the individual level, heuristic behavior is the most common mode of behavior with about 40% of subjects exhibiting it. Thus, the Diagnostic model is successful in demonstrating the robustness of the heuristic even in an environment in which traders make the simplest possible timing decisions.

In addition to explaining phenomena in real markets, the heuristic provides a potential explanation for herding observed in past sequential trading laboratory experiments (Cipriani and Guarino (2005) and Drehmann et al. (2005)). It suggests subjects herd as a function of their beliefs about the asset value, not due to conformity or beliefs about mistakes by others.[9] Because trade timing is exogenous in this literature, it is more difficult to identify the heuristic.[10]

An important issue in assessing the external validity of heuristic behavior is to determine its root cause. I demonstrate numerically that prospect theory preferences (Kahneman and Tversky (1979,1992)) are a promising candidate for explaining this behavior and argue informally that the heuristic is inconsistent with several other common behavioral explanations

[8]Rather than receiving a single signal as in Kendall (2015), subjects are always given a private signal in the period they arrive. They then decide whether or not to acquire an *additional* signal before trading. This modification enriches the strategy space because traders can condition their timing decision on the private information they obtain upon arrival. I show theoretically that traders with different initial signals must follow the same timing strategy in equilibrium. However, in the data, those whose signals confirm the price trend tend to trade earlier. Thus, endowing subjects with private information before they make their timing decisions allows for a better understanding of what drives their behavior.

[9]Cipriani and Guarino (2005) consider and reject the idea that herding is due to beliefs about mistakes by previous traders. Drehmann et al. (2005) consider and reject conformity. To date, no explanation of non-equilibrium herding that is supported by the data has been put forth in the literature.

[10]Park and Sgroi (2012) also relate herding to trade timing, but in an environment with three states and signals where *rational* herd behavior is predicted. While they do not have precise theoretical timing predictions, they argue (and observe) that those who have signals that are more likely to herd move later than those with monotonic signals. In contrast, subjects that herd move earlier here.

for experimental data.

The paper is organized as follows. In Section 2, I consider the main model. I derive theoretical predictions, outline the experimental treatments, and provide results. I then investigate non-equilibrium behavior in more detail, developing a heuristic to explain it. Section 3 develops and tests the Diagnostic model. Section 4 discusses potential implications of the discovered heuristic for actual markets and how it can potentially explain behavior observed in related experimental work. There, I also discuss how prospect theory generates behavior consistent with the heuristic. Section 5 concludes.

2 Model and Theory

2.1 Terminology

I distinguish between two related concepts: rushing and panics. *Rushing* refers to a trader who trades prior to receiving all possible private information. A *panic* occurs when multiple traders rush to trade simultaneously. In the main model, both rushing and panics are possible, whereas in the Diagnostic model, only rushing is possible because traders trade one at a time.

2.2 Model

Time is discrete with $t = 1 \ldots T$ trading periods. In each trading period, each of n risk-neutral trader may trade a single, common-value $V \in \{0, 1\}$, at a price established by a market maker (the experimentalist). The initial prior that the asset is worth $V = 1$ is p_1. When the asset value is realized at T, those who purchased the asset receive a payoff of $V - p$ and those who sold (short) receive a payoff of $p - V$. There is no discounting.

Each trader, identified by $i \in n$, receives a private signal before the first trading period, $\underline{s}_i \in \{0, 1\}$, which has a correct realization with probability $\underline{q} \in (\frac{1}{2}, 1)$. Each trader may trade only once in any of the T trading periods. If a trader waits until time T to trade, she receives an additional private signal, $\overline{s}_i \in \{0, 1\}$, immediately prior to T, which has a correct realization with probability $\overline{q} \in (\frac{1}{2}, 1]$. Note that I allow for the second private signal to reveal the true asset value perfectly.

Between each pair of trading periods, a binary *public* signal, $s_{P,t} \in \{0, 1\}$, is released that has a correct realization with probability $q_P \in (\frac{1}{2}, 1)$. A single price, $p_t = E[V|H_t] = Pr[V = 1|H_t]$, equal to the expected value of the asset conditional on all publicly available information, H_t, is set by a market maker prior to each trading period.[11] H_t includes all

[11]Incorporating a bid-ask spread would make interpretation of the data much more difficult because

prior trades, timing decisions, and prices and is observed by all traders.

This model may be thought of as a simplified version of the situation arising around a firm's earnings announcement. At the time of announcement, traders receive independent (low quality) private information about the revision in the firm's valuation. After the announcement, public information slowly becomes available through a series of public signals and, finally, some time later (perhaps at the filing of the firm's 10-K), more information about the firm's revised valuation becomes publicly known.[12,13]

2.3 Theoretical Predictions

Due to the richness of the strategy space, a complete characterization of all equilibria for general parameters is tedious to derive. I focus instead on the particular sets of parameter values used in the two treatments, R and W, provided in Table 1. The optimal trading strategies are provided in Proposition 1. All details and proofs are given in Online Appendix A.

Proposition 1: *In equilibrium:*
a) traders who trade prior to period T buy if $\underline{s}_i = 1$ and sell if $\underline{s}_i = 0$.
b) traders who trade in period T buy if the true asset value is 1 and sell if the true asset value is 0.

Proposition 1 states the standard result that traders optimally buy when their private belief is greater than the public belief (which is equal to the price), and sell otherwise. For a risk-neutral trader, this strategy leads to a positive expected profit, while doing the opposite would lead to an expected loss.

Proposition 2 provides a necessary condition for any equilibrium of treatment R and a statement of the unique equilibrium of treatment W.[14]

optimal timing strategies involve mixing. For a detailed analysis in a closely related model, see Kendall (2015).

[12] The slow release of public information, while not unreasonable, has the additional benefit of facilitating unique equilibrium predictions.

[13] Although earnings announcements are typically thought of as public information, private information can arise if traders have heterogeneous abilities to process information or have different models of a firm's dividend-generating process (Kandel and Pearson (1995)).

[14] Proposition 2 part a) does not specify the off-equilibrium timing strategies, but the proof ensures that all trades occur immediately independent of these strategies.

Proposition 2:
a) In any equilibrium with $\underline{q} = \frac{3}{4}$, $q_P = \frac{17}{24}$, *and* $\overline{q} = 1$ *(R), all trades occur in the first trading period.*
b) In the unique equilibrium with $\underline{q} = \frac{13}{24}$, $q_P = \frac{17}{24}$, *and* $\overline{q} = 1$ *(W), all trades occur in the final trading period.*
Equilibrium trading (buy or sell) strategies are given by Proposition 1.

Although the model is simple to describe, its analysis is complicated by the fact that the strategy space of a trader, even when restricted to a single trade, is very large. I look for an equilibrium in Markov strategies, but even then strategies still depend upon a player's initial signal, the price, and the number of traders that have traded prior to her. The proof of Proposition 2 first establishes that traders with different initial signals must follow the same strategy, simplifying the analysis greatly. I then establish a pair of sufficient conditions to ensure that traders trade at either $t = 1$ or $t = T$.[15]

The difference between R and W lies solely in the initial signal quality. In R, the better quality initial signal increases the profit from trading immediately. The proof of Proposition 2 establishes that, even if all other traders wait to trade until $t = T$, the price impacts of the public signals alone are sufficiently costly to ensure a trader must trade at $t = 1$. When combined with the relatively large price impacts of others' trades, one may anticipate a robust rational panic in which all traders rush to trade at $t = 1$.

In W, the initial signal quality is so poor that it rationally pays to wait to learn the asset value in the final period before trading. Even if all other traders trade before $t = T$, their trades produce so little price impact that it is still rational to wait. The equilibrium prediction aside, one may expect trends in prices to prove tempting for traders that use the intuition outlined in Section 1, given the many opportunities to trade before they receive additional information.

Proposition 2 provides point predictions about subject behavior, but subjects are unlikely to behave rationally 100% of the time. A weaker hypothesis is that subjects rush more often in R than in W. I formally state this comparative static result as Corollary 1 which follows immediately from Proposition 2.

Corollary 1: *Subjects trade earlier in R than W.*

[15]More generally, the game becomes a coordination game in which multiple equilibria may exist. It may be interesting to study how traders behave in this coordination game, but studying behavior under environments with unique equilibrium predictions seems like a logical first step.

2.4 Experimental Design

All subjects were recruited from the University of British Columbia student population using the experimental recruitment package Orsee.[16] Subjects came from a variety of majors and no subject participated in more than one session. I conducted four sessions of each treatment, for a total sample of 64 subjects ($n = 8$ subjects in each session). New randomizations were performed for each session's asset values and signals in order to avoid the possibility of a particular set of draws influencing the results. In each session, subjects first signed consent forms and then the instructions (provided in Online Appendix B) were read aloud. Subjects were allowed to ask questions while the instructions were read and then completed a short quiz. All quiz questions had to be answered correctly by each subject before the experiment began, and this policy was common knowledge. Once the experiment began, no communication of any kind between subjects was permitted.[17]

In each session, 30 paid trials, preceded by two practice trials, were run. In each trial, subjects made their trading decisions via computerized interfaces, an example of which is provided in Online Appendix B. Past prices and trades were available on an intuitive graphical display. Software was developed using the Redwood package (Pettit and Oprea (2013)) which uses HTML5 to allow for rapid updating of the computer interface. This feature allows for many more trials than would have been possible otherwise, which is important because learning in this relatively complex environment is seen to play a role (see Appendix A). An additional benefit of a large number of trials is that it provides a sufficient number of observations to study individual behavior, something which has not been fully explored in previous trading experiments.[18]

The trading environments were framed as such: subjects were told that they would trade a stock with the computer. It was emphasized that they could only trade once and that they must trade (in each trial).[19] Asset values were represented visually as bins containing different numbers of colored balls, with signals corresponding to draws from the appropriate bin.

Subjects earned a payoff in each trial as described in Section 3.1. The asset value, V, and prices were scaled by a factor of 100 currency units. Subjects were endowed with 100 units

[16]http://www.orsee.org/

[17]Subjects were separated by physical barriers so that they could not observe each others' private information or decisions.

[18]Cipriani and Guarino (2009) use a semi-strategy method in a related experiment as alternative way of obtaining many individual observations.

[19]Both Cipriani and Guarino (2005) and Drehmann et al. (2005) allow subjects to not trade in some treatments, finding that a considerable fraction do so. Given that not trading is never optimal and the additional complexity of the environment here, I chose to require subjects to trade in order to eliminate one potential source of noise.

with which to trade prior to each trial, for a maximum possible earning of 200 currency units per trial. In order to induce risk-neutrality, each currency unit represented a lottery ticket with a 1/200 chance to earn \$1.00 Canadian. After all trials were completed, a computerized lottery was conducted for each paid trial and subjects were paid according to the results of the lotteries. In addition, each subject was paid \$5.00 as a show-up fee. Average earnings were \$21.53 (minimum \$12.00, maximum \$30.00) with a corresponding wage rate over an hour and a half of \$14.35/hour.

Ex-ante assumptions about behavior must be made in order to set prices in the experiment. I assume that no information is revealed by the decision to wait, as is the case in equilibrium (see Online Appendix A). Therefore, if no trade occurs, the price is updated due to the information contained in the public signal only. After a trade, the price is updated according to Bayes' rule, assuming that traders follow equilibrium buy and sell strategies. In the case of a trade that occurs after an off-equilibrium timing decision, the price is set assuming traders make the optimal buy or sell decision according to their private information after the deviation.[20]

Subjects were told that prices reflect the mathematical expected value of the value of the asset, conditional on all public information. They were also explicitly told that prices would increase after buy decisions and favorable public signals, and conversely for sell decisions and unfavorable public signals. The exact amount of each increase or decrease was not communicated, but subjects participated in many trials so that they had an opportunity to learn about price movements over time.

2.5 Results

In all reported results, I focus on the last third of trials (10 trials). Because the trading environment is relatively complex, focusing on the latter third of trials eliminates some of the noise associated with early behavior as subjects have had time to learn about the environment. Evidence of this learning is provided in Appendix A.

To categorize trading directions, I formally define four possible trade behaviors.

[20] As shown in Section 2.5, this assumption is valid the vast majority of the time, but the assumption that a wait decision reveals no information is violated frequently in the data. In Section 4, I discuss the consequences of traders understanding that other traders may not be making optimal trading or timing decisions when prices are set assuming that they do.

Definition 1: *A subject's trade is classified as:*

1. *'Rational' if she trades in the direction indicated by her private information (Proposition 1).*
2. *'Herding' if she buys at a price greater than* 0.5 *with private information indicating that she should rationally sell, or sells at a price less than* 0.5 *with private information indicating that she should rationally buy.*
3. *'Contrarian' if she sells at a price greater than* 0.5 *with private information indicating that she should rationally buy, or buys at a price less than* 0.5 *with private information indicating that she should rationally sell.*
4. *'Irrational', if at a price of* 0.5*, she buys with private information indicating that she should rationally sell, or sells with private information indicating that she should rationally buy.*

Table 2 reports the trading results of the model. We observe that a relatively high percentage of behavior is rational which suggests that, at least along this dimension, traders have a good understanding of their environment. In treatment W, more herding trades are observed, which is the first indication that something other than rational behavior is present in this treatment.

To compare timing behavior across the R and W treatments, Figure 1 plots the empirical cumulative distribution functions (cdfs) of the fraction of trades that occur in periods $t \leq t'$ as a function of $t' \in 1 \ldots 8$. The rational prediction is that all trades occur at $t = 1$ in the R treatment and at $t = 8$ in the W treatment, corresponding to step functions that transition from 0 to 1 at $t' = 1$ and $t' = 8$, respectively. Figure 1 provides clear evidence that trades occur earlier in treatment R than in treatment W. The median trading periods in the four sessions of R and W are $\{1, 1, 1, 1\}$ and $\{2, 3, 3.5, 4\}$, respectively. Applying a non-parametric Mann-Whitney U test, one rejects the null of equal median trading periods (test statistic = 0, p-value = 0.05), supporting Corollary E1.[21] I capture these trading and timing results in the following two findings.

Finding 1 *(Proposition 1) In treatments R and W, approximately 85% of subjects rationally reveal their private information through their trades.*

Finding 2 *(Corollary 1) Subjects trade earlier in the treatment for which the equilibrium prediction is to trade at* $t = 1$ *(R) than in the treatment for which the prediction is to trade at* $t = 8$ *(W).*

[21] A Kolmogorov-Smirnov test also rejects the null of equal empirical cumulative distributions (p-value = 0.01).

Behavior in the R treatment provides a vivid demonstration of predictable, rational panics and their consequences. Traders trade in the first period almost 75% of the time, and over 90% of trades occur in the first two trading periods. Extreme clustering of trades is observed as subjects scramble to be the first to trade in all four sessions of the treatment. As first studied in Kendall (2015), rational rushing can produce informational losses due to traders acting on low quality information. In treatment R, when all traders rush to be the first to trade, each trades with information that is correct only 75% of the time, forgoing the opportunity to obtain perfect information. In fact, only one subject in one trial (out of 320 subject-trials) ever waits to obtain the asset value at $t = 8$, so that prices reflect the true asset value perfectly only 0.3% of the time.

Unlike behavior in the R treatment, behavior in the W treatment clearly departs from the rational prediction. Before digging deeper into this behavior in the following section, I first note an immediate consequence: informational losses are greater than predicted. Even though the asset value should be revealed for certain, trials exist in which none of the eight subjects waits to learn it and therefore prices differ from fundamental value after the game is over. I provide more formal evidence of informational losses produced by non-equilibrium behavior in Appendix B. Both too frequent rushing and trading against one's private information contribute to these additional losses. Finding 3 summarizes the results related to informational losses.

> **Finding 3:** *Rational panics result in considerable information losses, as predicted (R). Non-equilibrium panics and trade directions exacerbate these losses, even after subjects acquire substantial experience.*

2.5.1 Heuristic-Driven Behavior (Treatment W)

I now turn to an explanation for the non-equilibrium behavior observed in treatment W. As suggested as a possibility in Section 1, it appears that a large proportion of subjects make decisions based upon maximizing their chances of 'being right' (buying good assets and selling bad), rather than maximizing their expected profit. With private information, these criteria do not necessarily lead to the same trading decisions. To take a simple example, consider the case in which the price has risen to 0.9 but one's belief is 0.8 because one has unfavorable private information. Buying results in correctly purchasing a good asset 80% of the time, while selling results in correctly selling a bad asset only 20% of the time, so that one buys if maximizing the probability of being right. However, if one accounts for the sizes of the potential gains from trade, one should sell to maximize expected profit. Traders that want to maximize their chances of being right buy at beliefs greater than one half and sell otherwise.

Extending this intuition, if one wants to maximize the probability of being right, one will value *public* information that provides additional certainty about the asset value (even though it is fully reflected in prices). Therefore, one may be willing to wait to observe others' trades until sufficiently certain about the asset value, even if no additional private information is received. In treatment W, 62% of trading decisions occur in the intermediate periods, $t = 2, 3, \ldots, 6$, suggesting that subjects do in fact place value on public information.

To formalize a description of how subjects behave, I define what I refer to as the τ-herding heuristic. τ is the sufficiently high belief threshold beyond which a subject immediately buys. The heuristic is labeled with the term 'herding' because, as in the previous example, subjects that buy at high beliefs may herd, trading against their private information. In the following definition, and throughout the analysis of timing decisions, I assume subjects treat the asset values of 1 and 0 symmetrically.[22] Under this assumption, the price range can be transformed to $p' \in [0.5, 1]$ where $p' \equiv \max(p, 1-p)$, and similarly for beliefs. For convenience, I often to refer to prices and beliefs transformed in this manner simply as prices and beliefs.

Definition 2: *A trader uses the τ-herding heuristic if she has a threshold private belief, $\tau \in [0.5, 1)$, and uses the following strategy:*

1. *With a private belief, $b \in [1-\tau, \tau]$, always wait for more information, if possible. If not possible, trade according to private information.*
2. *With a private belief $b > \tau$, buy immediately. With a private belief, $b < 1-\tau$, sell immediately.*

Use of the heuristic, while non-optimal, implies that subjects are reacting to their environment in a sophisticated and predictable way such that formal tests of the behavior are possible. In developing the predictions, I make no assumption about the distribution of τ values in the population: qualitative predictions are possible without any such an assumption. On the other hand, I later show that it is possible to estimate τ on an individual basis.

If subjects follow the τ-herding heuristic, we expect to see trading at intermediate periods

[22] Given the neutral framing of their environment, there is no reason to think subjects would favor trading in one direction over the other. Formally, I assume symmetric behavior around $p = 0.5$: a trader facing a price, p, with private information, I, makes the same timing decision as a trader facing price, $1-p$, and complementary private information, I^C. To test this assumption, I partition trials into those in which $V = 1$ and those in which $V = 0$. If there were an asymmetry between rising and falling prices, one would expect the determinants of rushed trades to be different in these two samples. However, the interaction terms are insignificant if one interacts a dummy for $V = 1$ with the other covariates in the regression of Table 3 that follows.

when beliefs cross traders' threshold beliefs, as observed. To investigate these trades in more detail, consider what we expect to see if τ-herding subjects' threshold beliefs vary. No simple relationship between beliefs and the probability of trade at a particular price (or period) is predicted because of selection issues: traders with low threshold beliefs exit earlier in time. Instead, the prediction is that when beliefs exceed a particular level for the first time, we should observe all traders with threshold beliefs below that level trade, if they haven't already. I use the exogenous public signals, which drive most of the changes in prices (and therefore beliefs), as a coarse predictor of when belief thresholds are crossed. Specifically, I construct a set of indicator variables, each of which is set to one when the absolute value of the difference in public signals reaches a new value *for the first time* in the period prior to the trading period of interest (and is zero otherwise). I denote these dummy variables PubDiff1 through PubDiff6 (6 being the largest absolute difference in public signals arising in the data). One expects to see a higher probability of trade in the period immediately after a new level is reached, as long as there are traders with threshold beliefs between adjacent levels.

Table 3 reports the results of a logit regression of the probability of trade on the public signal difference indicator variables described above. I exclude $t = 8$ because traders must trade in this period if they haven't traded previously. I also control for subject and trial fixed effects, the price, whether a trader's signal makes their beliefs more or less extreme (i.e. a positive signal makes beliefs more extreme when $p \geq 0.5$, but makes beliefs less extreme when $p < 0.5$), and a linear time trend (i.e. trading period).[23] Results are also reported when the subsample is restricted to only subjects identified as τ-herding types, as discussed later in this section.

Table 3 shows that when public signals first push prices to about 0.7 (PubDiff1) and about 0.85 (PubDiff2), the probability of trade jumps, suggesting that subjects with τ values below ≈ 0.7 and between ≈ 0.7 and ≈ 0.85 both exist in the data.[24] The indicator variables corresponding to public signal differences of three or more are not significantly positive and one is in fact (marginally) significantly negative. But, because higher public signal differences are reached less often in the data, statistical power is low in these cases. For this reason, I cannot rule out the existence of τ-herding traders with threshold beliefs beyond 0.85.

The fact that a signal that makes one's belief more extreme does not increase the probability of trade is surprising, although it is positive as one would expect.[25] The positive

[23]The results are robust to clustering standard errors at the session level rather than using heteroskedastic robust standard errors. Because there are only a small number of clusters (4) and because the clustered standard errors are typically smaller, I choose to report robust standard errors instead.

[24]These characterizations are coarse because any trades by other traders, as well as one's own private information, also affect beliefs (to a smaller degree).

[25]Here, the effect of one's signal is averaged across cases in which a new public signal difference level is

coefficient on the time trend shows that traders are more likely to trade in later periods, all else equal.

I have provided timing evidence consistent with the heuristic's prediction, but it also predicts trade directions: subjects trade in the direction of the price trend. Because I have sufficient individual observations, I explore this dimension of behavior at the individual level, providing a more detailed picture than an aggregate analysis. I attempt to classify each individual into one of four possible types: rational, τ-herding, *τ-contrarian*, and *simplistic*. A trader is τ-contrarian if she follows a strategy identical to that of a τ-herding type except that she trades against her belief, rather than with it. In classifying an individual as τ-herding or τ-contrarian, I only require that their behavior be consistent with some value of $\tau < 1$. Specifically, a rushed trade may occur in any period prior to the final trading period, but must be in the direction prescribed by the heuristic.[26] Simplistic types always rush and follow their signal when they trade.

To classify subjects, I impose a rather stringent criterion: *both* their trading and timing decisions must be consistent with those of a particular type in nine of the last ten trials. I allow subjects to be classified into more than type in order to get a sense of the robustness of the classification, but I also assign a unique type based upon the following prioritization scheme (high to low): rational, simplistic, τ-herding, and τ-contrarian. Table 4 provides the resulting classifications for each treatment, including the percentages of exact matches and subjects that have ambiguous types. Given that types must match behavior along two dimensions, the classification scheme is quite successful: 84.4% of subjects are able to be classified.

About one third of subjects are classified as τ-herding types, which is more than five times as many than are classified as rational. Only the simplistic types are (slightly) more common, but a subjective evaluation of these types suggests that most are likely τ-herding types who made a single mistake too many to be classified as such. Supporting this assertion, in the Diagnostic treatments that follow, simplistic types mostly disappear such that across all treatments with substantial non-equilibrium behavior, τ-herding types are by far the most common.

I perform three additional checks that those individuals classified as τ-herding types actually follow the heuristic. First, the average belief at which these types trade should be

reached and when it is not. In Table 6, we see that when we consider its effect conditional on reaching the first threshold, it is significant.

[26] If one of these types trades at $t = T$, I require that their trade direction be consistent with either the price trend or their private information. As such, all rational subjects are also classified as τ-herding. For this reason, I prioritize rational types over τ-herding types in assigning the final classification. If I were to instead require τ-herding types to never wait, I would not be able to identify those that use relatively high belief thresholds.

higher than the average belief at which they wait. This criterion is satisfied for 81.8% of these subjects. Second, returning to the determinants of rushing in Table 3, for the eleven subjects that match τ-herding types we see that the increased probabilities of trade at the new public signal difference levels are more than double those in the aggregate regression. Although not significant, even the third public signal difference level has a large, positive coefficient. In this regression, a subject's signal now has a significant positive effect, indicating that traders with signals that make their beliefs more extreme are significantly more likely to trade, a finding I return to in Section 2.5.2. Lastly, following a procedure outlined in Appendix C, I estimate the individual belief thresholds,τ_i, for each of these individuals using only their timing decisions. 45.5% of them have thresholds significantly less than one, with thresholds ranging from 0.62 to 0.74. Most of the remaining subjects use higher thresholds that do not differ significantly from one, but some subjects appear to not use a stable threshold. Together these results suggest that a fairly heterogeneous mix of belief thresholds are used among those identified as τ-herding types. Finding 4 summarizes the evidence thus far on behavior.

> **Finding 4:** *Subjects exhibit considerable heterogeneity in their strategies. About one third of subjects can be described as τ-herding, which is more than five times the percentage of subjects that can be described as rational or τ-contrarian.*

One may wonder why heuristic behavior is completely absent in treatment R. In early trials of the R treatment, evidence in fact exists that subjects use the heuristic (results available upon request). But, the fact that its use causes large reductions in expected profits causes it to die out in later trials.[27] So, while most subjects naturally tend to follow the heuristic, if its use proves costly enough, they do learn their way out of it. Note, however, that they only learn their way out of it when panics are rational: little evidence exists that they learn to wait in treatment W even though one would think that learning the asset value perfectly provides a very salient reason to wait. Overall, we are left with panics, either rational or heuristic.

2.5.2 Implications of Heuristic Behavior

This section studies two implications of heuristic behavior: non-equilibrium panics (trade clustering) and positive serial correlations in returns.

[27] When all others are trading at $t = 1$, if one delays one's trade by even one period, the price (at $t = 2$) is likely to be very close to the true asset value, so that trading produces almost no gain over one's endowment of \$1.00 (100 currency units). Trading immediately instead results in an expected profit of $0.75 * \$1.50 + 0.25 * \$.50 = \$1.25$.

The increased probability of trade when new levels of public signal differences are reached provides indirect evidence of non-equilbrium panics, defined as clustered trades. As further evidence, consider the trading period ($t = 2$) after the first public signal is revealed (corresponding to PubDiff1), where the increase in trade probability is the largest. $81/320 = 25.3\%$ of all trades occur in this second trading period even though it makes up only 12.5% of all trading periods. Put another way, there are an average of 2.03 trades at $t = 2$. If the eight traders were randomly choosing to trade in one of the eight periods, the expected number of trades at $t = 2$ would be given by a binomial distribution, so that we'd expect only $8 * 0.125 = 1$ trade, on average. Under the binomial distribution, the probability of observing more than 2.03 trades is only 6.25%. Thus, we see statistical evidence that non-equilibrium panics occur in a predictable manner.

When τ-herding subjects rush to trade, they trade in the direction of their private belief, producing positive correlation in returns. To study this correlation, I continue to focus on the second trading period just after the first public signal is revealed. At this time, 85% of trades are in the direction of the first public signal (and therefore in the direction of the trader's belief). Table 5 provides the Spearman correlation coefficients between the return due to the first public signal and that due to trades in each of the first seven trading periods, excluding the last where the asset value is known.[28] We observe a very strong positive short-term correlation between the return due to the first public signal and that due to the trades in the following trading period. Were traders simply rushing and trading according to their signals, prices would form a martingale and such a correlation would be statistically improbable. This finding provides solid evidence of trades being not only clustered, but also occurring in the same direction.[29] Finding 5 summarizes the clustering and correlation results.

Finding 5: *Use of the τ-herding heuristic produces non-equilibrium panics (clustered trades) and positive short-term correlation in returns.*

I investigate further the source of the correlation at $t = 2$, providing two pieces of evidence to show that it is primarily driven by traders with different private information choosing to trade at different times, and *not* by herding. Although herding has long been considered to be a source of correlation (Hirshleifer and Teoh (2003)), this finding suggests a different explanation.

[28]The results are very similar if the Pearson correlation coefficient is used, but there is no a priori reason to expect the correlations to be linear.

[29]One may expect negative correlation in later trading periods, should those with signals that did not agree with the first public signal delay their trades and later trade in the opposite direction. However, any such negative correlation would be spread over the remaining trading periods as there is no salient time at which to trade. In addition, some of these traders may actually delay until $t = 8$ when they learn the true asset value.

First, of the 81 trades that occur at $t = 2$, 65% are by those whose signal agrees with the public signal even though only 51% have such signals. This fact is consistent with τ-herding behavior because traders whose signals agree with the public signal have more extreme beliefs so that their critical threshold for trading is more likely to be exceeded. Furthermore, 100% of those whose signals agree with the public signal trade in the direction of their beliefs, while only 57% of the others do (i.e. 57% herd). Therefore, on average, the herding trades contribute little to the price change.

As a second piece of evidence, I show that the increased probabilities of trade at new public signal levels are driven by those whose signals make their beliefs more extreme than the current price. To do so, I interact each of the public signal indicator variables in Table 3 with a dummy that indicates a trader's signal makes her belief more extreme. The results of Table 6, show that the increased probability of trade after the first public signal is solely due to traders whose signals make their beliefs more extreme.[30] The coefficient on the first public signal indicator, which now represents those with signals whose beliefs are less extreme than the price, is a fairly precise zero, indicating that there is no increased probability of a trade that could result in herding.

In summary, we have convincing evidence that correlation in returns arises because traders choose when to trade based upon their private information. This finding provides a new explanation for the empirical puzzle of correlation in stock returns, specifically post-earnings-announcement drift, in which returns drift in the direction of the earnings surprise (here, the first public signal).[31]

3 Diagnostic Model and Theory

In the main model, the strategy space of subjects is very large, so that one may conjecture that complexity leads to non-equilibrium behavior. Also, selection issues complicate the heuristic predictions. For these reasons, I also study a simpler model in which traders make binary timing decisions and trade one at a time. Although panics are precluded in the model, it serves as a useful diagnostic tool.

3.1 Diagnostic Model

The model is based on that of Kendall (2015). The asset structure, common prior, and payoffs are identical to the main model. In each period $t = 1, 2, \ldots, T$, a single new, risk-neutral

[30] For brevity, I do not report the coefficients on all variables. The price coefficient remains insignificant and the period coefficient remains positive and significant.

[31] See Daniel et al. (1998) for a list of papers providing evidence of post-earnings-announcement drift.

trader arrives to the market and may trade the asset of unknown value.

Upon arrival, each trader receives a private, binary signal, $\underline{s}_t \in \{0, 1\}$, which has a correct realization with probability $\underline{q} = Pr(\underline{s}_t = 1|V = 1) = Pr(\underline{s}_t = 0|V = 0) \in (\frac{1}{2}, 1)$. The trader may either buy or sell immediately in the period she arrives or wait (not trade) and trade in the following period. Each trader may trade only once. If a trader waits, she receives an additional private signal before she trades , $\overline{s}_t \in \{0, 1\}$, which has a correct realization with probability $\overline{q} = Pr(\overline{s}_t = 1|V = 1) = Pr(\overline{s}_t = 0|V = 0) \in (\frac{1}{2}, 1)$. All signals are independent conditional on V.

Should a trader choose to wait, her next trading opportunity is immediately *after* the next trader arrives to the market, as shown in Figure 2. As indicated by the highlighted (second) trader in Figure 2, all traders (except the first and last) face up to two intervening trades: one from trader $t-1$ (if she chose to wait) and one from trader $t+1$ (if she chooses to rush). The complete history of trades, timing decisions, and prices (denoted $\underline{H}_t$ in the subperiod of a new arrival and $\overline{H}_t$ in the subperiod reached after a trader waits) are observed by all traders. Traders may buy or sell at a single price equal to the expected value based upon all public information, $\underline{p}_t = E[V|\underline{H}_t] = Pr[V = 1|\underline{H}_t]$ or $\overline{p}_t = E[V|\overline{H}_t] = Pr[V = 1|\overline{H}_t]$. $\underline{p}_t$ plays a more prominent role, so I generally refer to at as 'the' price and denote it p_t, unless a distinction is necessary.

In designing the two treatments for the Diagnostic model, Diagnostic Rush (DR) and Diagnostic Wait (DW), the goal is to provide a strong comparative static test across treatments: subjects should rush in one, but wait in the other. To achieve this, I chose the parameters for the two treatments specified in Table 7. For both treatments, the initial prior, p_1, is set to $\frac{1}{2}$. The two treatments differ mainly in the quality of the initial information received, being higher in Diagnostc Rush than Diagnostic Wait. I set $T = 6$ so that there are 6 subjects in each session of each treatment.

The experimental protocol for the Diagnostic treatments is virtually identical to that of the main treatments discussed in Section 2.4. Four sessions of each treatment, for a total sample of 48 subjects, were run. Each session consisted of 42 trials plus two practice trials. Subject orderings were randomized in each trial. Average earnings were \$28.67 (minimum \$22.00, maximum \$35.00) with a corresponding wage rate over an hour and a half of \$19.11/hour. The instructions and an example of the trading interface are provided in Online Appendix B. Prices are set as in the main model. If no trade occurs, the price remains unchanged.

3.2 Theoretical Predictions: Diagnostic Model

For general signal qualities, $\underline{q}$ and $\overline{q}$, traders in the model face a non-trivial trade-off between the value of additional information and the potential cost of adverse price movements. Kendall (2015) establishes a full equilibrium characterization for a closely related model, but here I consider only the parameterizations used in the experiment.[32] I present the results and intuition, leaving the details to Online Appendix A.

The optimal trading strategies for the Diagnostic model are provided in Proposition D1.

Proposition D1: *In the Diagnostic model:*
a) traders who rush in either treatment and traders who wait in the DR treatment buy if $\underline{s}_t = 1$ *and sell if* $\underline{s}_t = 0$.
b) traders who wait in the DW treatment buy if $\overline{s}_t = 1$ *and sell if* $\overline{s}_t = 0$.

As in the main model, the trading strategies of Proposition D1 are optimal because they ensure one buys when one's private belief is greater than the public belief (which is equal to the price), and sells otherwise. If one has contradictory signals, one trades according to the one with higher quality.[33]

Proposition D2 specifies the equilibria of the Diagnostic model for the parameterizations used in the experiment.

Proposition D2:
a) In any equilibrium of the Diagnostic model with $\underline{q} = \frac{17}{24}$ *and* $\overline{q} = \frac{16}{24}$ *(DR) all traders rush at every history, except trader* T*, who is indifferent when trader* $T-1$ *rushes.*
b) In the unique equilibrium of the Diagnostic model with $\underline{q} = \frac{13}{24}$ *and* $\overline{q} = \frac{17}{24}$ *(DW) all traders wait at every history.*
Equilibrium trading strategies are given by Proposition D1.

In the DR treatment, rushing is intuitively optimal because the second signal never changes one's trading decision: if the two signals contradict one another, trades are made according to the first, better quality, signal. Thus, there is no benefit from waiting, but there is a cost due to others' trades. The final trader is an exception in that she is indifferent between rushing and waiting when $T-1$ rushed, due to the fact that no intervening trades are possible. In the DW treatment, a strictly positive benefit from waiting exists for all

[32] The Diagnostic model is a non-trivial modification of Kendall (2015) in that, because traders receive private signals prior to choosing to rush or wait, they may choose to wait with different probabilities according to their 'type', defined as the private signal they receive upon arrival. As in the main model, both types must wait with the same probability in equilibrium (see Appendix A).

[33] To ensure traders are never indifferent, the second period signal strength is set slightly lower than the first period signal strength in treatment DR.

p_t, independent of whether or not $t-1$ and $t+1$ rush or wait, making the optimal timing strategy to wait. Here, the much better quality information obtained by waiting dominates the cost of any potential trades.

As in the main model, a weaker comparative static prediction follows directly from Proposition D2.

Corollary D1: *Subjects rush more often in DR than in DW.*

In contrast to the equilibrium predictions of Proposition D2, a simple heuristic prediction is available in this environment. If subjects follow the τ-herding heuristic and, as found in the main results, use heterogeneous threshold beliefs, then, in the aggregate, the decision to rush or wait should increase with a subject's belief upon arrival. In having subjects make only a single timing decision, selection issues are avoided. This prediction is independent of the treatment, but as observed in the main results, we may expect the heuristic to be moderated by the cost of its use.

3.3 Results: Diagnostic Model

As with the main results, I report data from only the last third of trials (14 trials) in order to focus on behavior after subjects have gained experience.

Table 8 first demonstrates that trading behavior in the Diagnostic model is highly rational, as in the main model. Finding 6 summarizes the trading behavior.[34]

Finding 6 *(Proposition D1) In the Diagnostic treatments, almost 85% of subjects' trade directions are rational, revealing their private information.*

Turning to the timing decisions, Table 9 provides the percentage of rush decisions observed in each treatment, aggregated across sessions. In the DR treatment, I omit the timing observations of the final trader when the previous trader rushed, leaving 309 observations.[35] A t-test comparing the average frequency of rushing across treatment rejects the null of equal means at the 5% level (p-value = 0.046), supporting Corollary D1. Subjects respond to the difference in equilibrium forces across treatments in a predictable manner. Interestingly, behavior is not statistically different across treatments in the first third of trials (difference in means is 9.0%, p-value = 0.24), indicating that subjects learn to avoid adverse prices move-

[34] One might expect a lower degree of rational behavior when a subject's two signals conflict. In Diagnostic Rush, this conjecture is correct, with subjects acting rationally only 60% of the time when they receive conflicting signals. In Diagnostic Wait, subjects rationally follow the better quality signal 90% of the time.

[35] When the final trader knows the previous trader rushed, she is indifferent. Of the 27 times traders are indifferent, they wait 66.7% of the time.

ments as they acquire experience. I summarize this comparative static result in Finding 7.

Finding 7 *(Corollary D1) In the Diagnostic treatments, subjects rush more often in the treatment in which it is the equilibrium prediction (Diagnostic Rush).*

Although subjects respond as predicted by Corollary D1, behavior is clearly not perfectly rational. In particular, subjects both wait too much (Diagnostic Rush) and rush too much (Diagnostic Wait) which is suggestive of the heuristic prediction that subjects wait at prices near one half, but rush as prices become more extreme.

To provide an overall view of the aggregate timing decisions, I plot the proportion of rushed trades as a function of price (Figure 3). I bin the data according to the number of trades that reveal better quality information (late trades in DW and any trade in DR) because these trades are the main determinants of prices. I also provide 95% confidence intervals, assuming that the decision to rush is binomially distributed. Note the clear upward trend in treatment DW: the probability of rushing increases with the price. Such a trend occurs in treatment DR as well, although it is notably noisier. These upward trends are not predicted by rational behavior and provide a first indication of heuristic behavior.

Table 10 presents the results of a logit regression of the probability of a rushed trade on a trader's private belief for each treatment. In the regressions, subject and trial fixed effects are included (not reported), as is an indicator dummy that indicates whether or not the previous trader has rushed.[36] Results are also shown for the subsample of individuals classified as τ-herding types later in this section.

In Table 10, we see that, consistent with the heuristic prediction, the probability of a rushed trade increases with a trader's belief. This effect is statistically and economically significant in the DW treatment: moving from complete uncertainty (belief is 0.5) to completely certainty (belief is 1), the probability of rushing increases by almost 0.85, a very large effect. In the DR treatment, we also observe a positive coefficient on a trader's belief but it is not statistically significant. Table 10 also shows that the probability of rushing, although not significant, is reduced when the previous trader rushed, as is expected because there can be no price impact from the previous trader in this case.

A good reason exists for the fact that behavior in the DR treatment is not as well captured by the heuristic: using the heuristic is more costly in this treatment. The cost from deviating from the rational strategy is largest at prices near one half where trading profits are the largest. Conversely, at prices near certainty, payoffs are flatter because trading profits are

[36] I again choose to report heteroskedastic robust standard errors rather than standard errors clustered at the session level. The significance of results remains unchanged with clustered standard errors.

smaller. Thus, in treatment DR, following the heuristic by waiting at prices near one half is very costly: strong payoff incentives and the heuristic oppose each other. Conversely, in DW, at prices near one half, the heuristic and payoff incentives reinforce each other because both induce one to wait. One would then expect behavior to be more noisy in treatment DR than in treatment DW, as we observe. The fact that the heuristic is moderated or strengthened by the cost of its implementation here and in the main results, suggests that it is not itself driven by changes in payoff incentives but is instead an innate tendency that subjects bring to the lab.

As in the main results, I classify subjects in each of the Diagnostic treatments into one of four types.[37] In the resulting classifications, τ-herding types are the most common type in both treatments, making up 40.0% of the population overall and 67.9% of subjects that can be classified. Rational types are second most common making up only 12.5% of the population. Simplistic types are almost non-existent. Demonstrating the robustness of the heuristic, τ-herding types are frequent in both treatments, even in DR where payoff incentives work against the heuristic. The classification scheme works much better in DW, where 79.2% can be classified. This fact is again consistent with heuristic behavior and payoff incentives reinforcing each other to produce less noise in the DW treatment.

As I did in the W treatment, I perform three additional checks to ensure that τ-herding types' behavior is consistent with the heuristic. First, if a subject uses a consistent threshold, τ, then the beliefs at which they rush to trade should be higher on average than the beliefs at which they wait. This criterion is satisfied for 84.2% of the τ-herding types. Second, I return to the regression results in Table 10 for the subsample of subjects classified as τ-herding. The previous relationship between the probability of rushing and one's belief is now stronger and significant in treatment DR, as expected. Finally, for those subjects classified as τ-herding types, I estimate their individual threshold, τ_i. Of these types, 42.1% have thresholds estimated to be significantly less than one. Even among the thresholds that are significant, substantial heterogeneity exists, with point estimates of τ_i between 0.65 and 0.89. Given the overall evidence that subjects use the τ-herding heuristic, we have Finding 8.

Finding 8: *Across all three treatments in which non-equilibrium behavior is observed, 37.5% of subjects use the τ-herding heuristic, making it the most frequent mode of behavior. By comparison, only 10.0% of subjects are classified as rational in these three treatments.*

[37] As in the W treatment, simplistic types in DW always rush and trade according to their initial signal. Simplistic types in DR always wait and trade according to their private information.

4 Discussion

A critical question facing any laboratory study of markets is the degree to which behavior exhibited by university students can be extrapolated to real-world financial settings with experienced professionals. In terms of rational behavior, Kendall (2015) cites a variety of empirical evidence that is consistent with the fact that traders rationally rush to trade. Data from the laboratory backs up this evidence by providing a clear demonstration of rational panics in the R treatment, where subjects clearly understand that waiting to trade can be extremely costly. Although individually rational, the resulting rational panic is severely costly for the market as a whole due to large informational losses. It also produces acute trade clustering, which, interpreted in terms of an earnings announcement, is consistent with the fact that volumes spike around these announcements (Frazzini and Lamont (2006)).

The laboratory data can also provide insight as to when rational panics are more likely to occur in actual financial markets. Comparing observed behavior across treatments in both environments, we should expect to see rational panics when there are smaller differences in quality between initial private information and the information that can be acquired through research. Also, panics are more likely when many other informed traders are present in the market (treatment R) than when there are relatively few (treatment DR). Finally, note that experience in the environments where rushing is optimal drives behavior *towards* rational panics. Given that equilibrium behavior involves preempting other traders, it seems likely that experienced professionals understand the benefit of acting quickly, and do so.

Although heuristics found in the laboratory should generally be cautiously exported to the field, there are good reasons to suspect the τ-herding heuristic operates in real-world markets. First, as already noted, the heuristic, although not optimal, is far from irrational behavior: it demonstrates a reasonable level of understanding of one's environment. Second, the percentage of τ-herding types in the data actually *increases* over time. From the first third of trials to the last third, the changes are from 12.5% to 34.3%, 25.0% to 50.0%, and 12.5% to 29.1%, in treatments W, DW, and DR, respectively. As in the case of rational panics, this finding suggests that traders may actually learn over time that prices move against them in expectation, causing them to (rationally or heuristically) panic. Finally, the fact that such behavior generates the well-known phenomena of short-term positive correlation in returns and trade clustering in the laboratory data suggests that it is a good candidate explanation for these findings.[38]

[38] Although not detectable in the laboratory data, given that traders with signals that oppose the price trend delay their trades, it is not inconceivable that τ-herding behavior could also be responsible for the longer-term negative correlation in returns (overreaction) observed in real markets (see Appendix A of Daniel et al. (1998) for a review of papers finding such evidence).

As further suggestive evidence that traders may exhibit τ-herding behavior in actual markets, consider the fact that one of the pillars of technical analysis is momentum trading: wait until prices appear to be moving in a particular direction and then trade in that direction.[39] In the absence of private information, momentum trading and τ-herding are identical, so, at least among those that subscribe to technical analysis, something akin to τ-herding strategies is natural.[40] Furthermore, τ-herding strategies can potentially generate positive-feedback trading (De Long et al. (1990)). As initial traders cause a trend in prices, those with more extreme threshold beliefs are encouraged to join in, reinforcing the trend and setting off further trading.

Uncovering the τ-herding heuristic by endogenizing trade timing also suggests an explanation for behavior in past experiments in which use of the heuristic was not easily identifiable. Cipriani and Guarino (2005) and Drehmann et al. (2005) study trading behavior in an exogenous timing version of the Diagnostic model. Both papers document herding and contrarian trades, but neither puts forth an explanation backed by the data. Cipriani and Guarino (2005) consider the possibility that subjects believe previous subjects made mistakes, but show that this explanation can only explain contrarian trades.[41] Drehmann et al. (2005) instead consider, but reject, conformity: subjects trade in the direction of previous subjects, regardless of their signals or beliefs.[42] The τ-herding heuristic suggests an explanation for herding: subjects herd when their beliefs become extreme.

Previous literature (De Long et al. (1990) and Cipriani and Guarino (2005)) points out that contrarian trades tend to stabilize markets by causing prices to become less extreme, while herding trades do the opposite. Thus, the finding that τ-herding types are much more frequent than τ-contrarian types suggests that behavioral types may tend to destabilize markets overall. This conclusion stands in contrast to Cipriani and Guarino (2005) who, based

[39] For a description of momentum trading, see http://www.investopedia.com/articles/trading/02/090302.asp.

[40] Note, however, that there is a subtle difference between τ-herding strategies and momentum strategies. In a momentum trading strategy, traders that observe a positive price change (or series of positive price changes) purchase the stock regardless of their private information, so there is no predicted relationship between traders' signals and when they choose to trade. Instead, through the use of a threshold belief, τ-herding behavior allows for differences in trade timing precisely because traders with different signals have different beliefs. As shown in Section 2.5.1, the correlation in returns in the data is primarily driven by differences in timing strategies and not by herding. If momentum strategies were instead the correct explanation, herding would be a more significant contributor to return correlation in the laboratory data. Several behavioral finance papers posit that traders follow momentum strategies in order to explain post-earnings-announcement drift (see De Long et al. (1990) and Hong and Stein (1999)), but the τ-herding behavior observed here provides an alternative explanation.

[41] I similarly reject this explanation as a source of herding in the data here (see footnote 45).

[42] I also reject conformity by looking for evidence in the W treatment where τ-herding and conformity have distinct predictions because past trades have relatively little effect on one's beliefs. Of the 51 herding decisions observed in this treatment, 42 coincide with new public signal difference levels, meaning traders generally herd *at the same time*, not after others have traded in a particular direction.

upon a higher frequency of contrarian trades in their data (Drehmann et al. (2005) find the same), conclude that contrarianism is more frequent. I suggest two reasons previous data may have led to a different conclusion. First, in the data here, the overall frequency of contrarian trades declines over time, from 13.1% in the first third of trials to only 5.3% in the last third. As these two previous studies did not allow for as many learning opportunities, they picked up behavior only in non-experienced subjects. Second, part of the reason contrarian trades may be more frequent is that there are more opportunities to act contrarian: signals are more likely to go in the direction of the price trend rather than against it. Conditioning on the opportunity for herding or contrarian behavior, I find much higher rates of herding behavior than contrarian behavior (47% vs. 12% in the main treatments and 21% vs. 14% in the Diagnostic treatments).

Given the explanatory power of the heuristic, two additional questions naturally arise. Do subjects understand that others use the heuristic? And, what is the underlying behavioral force that generates its use? On the first question, the data is not consistent with subjects understanding that others use the heuristic. First, the parameters in all treatments are such that one's timing decision is a dominant strategy so that following the heuristic's timing strategy is not optimal even if others do.[43] Second, prices are set in the experiment assuming that observing a trader wait reveals no information and that traders always trade according to their private information. τ-herding behavior causes violations of both of these assumptions so that prices are seen to be too extreme ex post.[44] If traders understand this fact and the mispricing is severe enough, then they may be induced to act contrarian and sell regardless of their signal (perhaps rushing to do so). Because such contrarian behavior is very infrequent, it would seem either that the majority of traders do not understand others are τ-herding types, or that the drive to follow the τ-herding strategy is so strong that traders follow it even after accounting for mispricing.

Can τ-herding behavior be generated by common behavioral game theoretical models or alternative preferences? Unfortunately, given the complexity of the models considered here, a full equilibrium characterization for any particular theory is extremely difficult. However, I argue that prospect theory offers the best chance of being at the root of the behavior. Most common behavioral theoretic models build upon the idea that subjects make mistakes, but if subjects make mistakes, then prices are generally too extreme. As argued above, subjects that account for mistakes, would then trade in a contrarian manner, but would never herd.[45]

[43]In treatment DR, one is indifferent if one's predecessor rushed and one expects one's successor to wait, but there is never a strict incentive to wait.

[44]For example, consider the case in which the price trend is upwards. A decision to wait reveals that a trader is more likely to have had a negative signal. Furthermore, a buy decision may be taken by a trader with a negative signal. In both cases, the resulting price is too high.

[45]Behavioral game theory has developed several theories that have proven successful at describing ex-

Turning to preference-based explanations, I first note that risk aversion is incapable of producing heuristic behavior. Intuitively, risk-averse individuals prefer to wait to learn about the asset value, potentially even through public information, which is consistent with one aspect of the heuristic.[46] However, it can not generate rushing when the risk-neutral prediction is to wait, nor buying with unfavorable information, and therefore fails to produce the other two aspects.

The prospect theory preferences of Kahneman and Tversky (1979,1992), on the other hand, can in fact produce τ-herding behavior if one is willing to accept that possibility that subjects exhibit preferences different from those induced (risk-neutrality was induced by paying subjects in lottery tickets). In fact, prospect theory's S-shaped utility function can produce all three characteristics of the heuristic: waiting at uncertain prices, rushing at more certain prices, and trading in the direction of the price trend. To demonstrate this possibility, I numerically simulated subject's timing and trading decisions in both trading models, making several simplifying assumptions.[47] First, I assume that other subjects follow rational strategies. Although clearly not the case in the data, an equilibrium characterization would be very difficult. In treatment W at least, this assumption is not critical because prices are mainly determined by the public signals. Second, I assume subjects evaluate gains and losses relative to a reference point equal to their endowment. Other reference points, such as one's expected profits (Koszegi and Rabin (2006)), are possible, but the endowment reference point is simple and seems natural in this environment. Third, I assume subjects'

perimental behavior: quantal response equilibrium (McKelvey and Palfrey (1995,1998)), level k reasoning (Stahl and Wilson (1994) and Nagel (1995)), cognitive hierarchy (Camerer et al. (2004)), and cursed equilibrium (Eyster and Rabin (2005)). These theories are based upon models of subjects' beliefs about other subjects behavior. Intuitively, they all lead to contrarian behavior. Quantal response equilibrium assumes that subjects probabilistically take actions in proportion to the payoff of each action, and that this behavior is common knowledge. While there is some evidence that mistakes are proportional to payoff differences (see the discussion in Section 3.3), if subjects believe others make such mistakes, then they would believe prices are too extreme. In this case, contrarian behavior would be prescribed. In level k and cognitive hierarchy, subjects believe other subjects are boundedly rational and best respond to this behavior. The lowest level of rationality is typically assumed to be random behavior, but this would again lead to prices being too extreme and contrarian behavior. Finally, in cursed equilibrium, subjects correctly predict the frequencies of actions, but believe that actions are not connected to underlying signals. Again, this belief leads to believing prices are too extreme and contrarian behavior.

[46]Assuming power utility, numerical simulations suggest though that no value of risk aversion can produce non-equilibrium waiting at prices near one half in treatment DR where it is observed.

[47]I also included the probability weighting and loss aversion aspects of prospect theory in the numerical simulations, but omit them from the exposition. Loss aversion appears to do little to change behavior. Probability weighting actually works against the heuristic, so that if a subject distorts probabilities, they must have more pronounced utility curvature to exhibit heuristic behavior. Probability weighting further complicates the analysis because it induces time-inconsistencies in behavior (see Barberis (2012). In fact, the main model is very similar to the casino gambling model of Barberis (2012), and so is computationally intensive to solve when there are many trading periods, as he shows.

utility functions are S-shaped and governed by power utility

$$u(x) = \begin{cases} (x-1)^\alpha & x > 1 \\ -(1-x)^\alpha & x < 1 \end{cases}$$

where the value 1 is a subject's normalized endowment.

For each of the three environments in which heuristic behavior is observed, numerical simulations show that for α sufficiently small, heuristic behavior is produced. To understand the intuition, consider an individual with no initial private information facing a high price, $p > \frac{1}{2}$. Buying provides a lottery in which one obtains a gain of $1-p$ with probability p, and a loss of p with probability $1-p$. Selling instead provides a gain of p with probability $1-p$ and a loss of p with probability $1-p$. Compare first the possible gains. Due to risk aversion over gains, obtaining a smaller amount, $1-p$, with a larger probability, p, is preferred to the converse, even though the expectation is the same. Due to risk-seeking over losses, the converse is true: one prefers the smaller probability, $1-p$, of a larger loss, p. Therefore, both in terms of gains and losses, a subject with prospect theory preferences strictly prefers to buy at any $p > \frac{1}{2}$, while a risk-neutral subject is indifferent. Similarly, at $p < \frac{1}{2}$, a subject with prospect theory preferences strictly prefers to sell. To see that prospect theory preferences induce waiting at uncertain prices, note that the utility of such a subject is zero at $p = \frac{1}{2}$, but strictly positive away from $p = \frac{1}{2}$. Therefore, a subject initially waits to allow others to trade, buying if they buy, and selling if they sell, just as the heuristic prescribes. Finally, for prices sufficiently close to 0 or 1, there is no longer an incentive to wait because one knows one will buy (or sell) regardless of future price movements. Adding in private information complicates the analysis, but the basic intuition remains unchanged.

Although prospect theory is a promising candidate to explain heuristic behavior, a caveat is in order. Numerical simulations show that the utility curvature must be very pronounced to produce herding even at extreme prices, and no amount of curvature can produce herding at the more intermediate prices at which it is observed in the experiment. Therefore, prospect theory as an explanation remains speculative. Going forward, it is natural to study it in greater detail in the simpler sequential trading model in which herding and contrarianism have been observed (Cipriani and Guarino (2005) and Drehmann et. al. (2005)). I am currently pursuing this line of research.

5 Conclusion

This paper investigates the sources and consequences of trading panics in a laboratory setting in which precise theoretical predictions are derived. Traders may rationally rush to trade in order to avoid adverse price movements that reduce their profits, leaving substantial information on the table. Theoretical results successfully predict the relative frequency of rational panics and allow non-equilibrium panics to be identified. Non-equilibrium panics result from a simple heuristic and exacerbate informational losses. The heuristic reconciles findings in previous experimental studies with related environments and puts forth a new explanation for the trade clustering and positive short-term correlation in returns observed in real financial markets.

A clear direction for future experimental research is to understand in more detail the trading heuristic used by almost half of traders. Of particular interest may be to determine whether or not it is truly generated by prospect theory preferences, and whether or not subjects that try to exploit such behavior to their advantage exist. For example, in settings in which multiple trades are possible, rational traders may buy in an attempt to induce heuristic traders to buy. If successful, they can then profitably sell at a higher price. In this way, the heuristic may help to explain price bubbles, both in the field and in laboratory experiments (e.g. Smith et al. (1988)).

The trading heuristic has additional implications outside of the laboratory. In particular, it suggests that short-term correlation in returns should be more dramatic after earnings announcements that generate disagreement among traders, as opposed to those that are unambiguous in their implications. In the former case, there is more scope for traders with different beliefs to choose to trade at different times. Similarly, return correlations should be more pronounced in stocks with higher informed trader arrival rates, where increased fears of adverse price movements cause traders to rush to trade on lower quality information. Given these predictions, future empirical studies may be able to detect the use of the trading heuristic even without being able to observe traders' private information.

References

[1] Avery, C. and Zemsky, P. (1998), "Multidimensional uncertainty and herd behavior in financial markets", *American Economic Review*, 88 (4), 724-748

[2] Banerjee, A. (1992), "A simple model of herd behavior", *The Quarterly Journal of Economics*, 107 (3), 797-817

[3] Barberis, N. (2012), “A model of casino gambling”, *Management Science*, 58 (1), 35-51

[4] Barberis, N., Shleifer, A, and Vishny, R. (1998), “A model of investor sentiment”, *Journal of Financial Economics*, 49, 307-343

[5] Barlevy, G. and Veronesi, P. (2003), “Rational panics and stock market crashes”, *Journal of Economic Theory*, 110, 234-263

[6] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992), “A theory of fads, fashion, custom, and cultural change as informational cascades”, *The Journal of Political Economy*, 100 (5), 992-1026

[7] Bloomfield, R., O’Hara, M., and Saar, G. (2005), “The ‘make or take’ decision in an electronic market: evidence in the evolution of liquidity”, *Journal of Financial Economics*, 75, 165-199

[8] Brunnermeier, M. and Morgan, J. (2010), “Clock games: theory and experiments”, *Games and Economic Behavior*, 68 (2), 532-550

[9] Brunnermeier, M. and Pedersen, L. (2005), “Predatory trading”, *The Journal of Finance*, 60 (4), 1825-1863

[10] Bulow, J. and Klemperer, P. (1994), “Rational frenzies and crashes”, *Journal of Political Economy*, 102 (1), 1-23

[11] Camerer, C., Ho, T, and Chong, J. (2004), “A cognitive hierarchy model of games”, *The Quarterly Journal of Economics*, 119 (3), 861-898

[12] Çelen, B. and Hyndham, K. (2012), “An experiment of social learning with endogenous timing”, *Review of Economic Design*, 16 (2-3), 251-268

[13] Chamley, C. and Gale, D. (1994), “Information revelation and strategic delay in a model of investment”, *Econometrica*, 62 (5), 1065-1085

[14] Cipriani, M. and Guarino, A. (2005), “Herd behavior in a laboratory financial market”, *The American Economic Review*, 95 (5), 1427-1443

[15] Cipriani, M. and Guarino, A. (2009), “Herd behavior in financial markets: an experiment with financial market professionals”, *Journal of the European Economic Association*, 7 (1), 206-233

[16] Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (1998), “Investor psychology and security market under- and overreactions”, *The Journal of Finance*, LIII (6), 1839-1885

[17] De Long, J.B., Shleifer, A., Summers, L., and Waldmann, R. (1990), "Positive feedback investment strategies and destabilizing rational speculation", *The Journal of Finance* 45, 375–395

[18] Devenow, A. and Welch, I. (1996), "Rational herding in financial economics", *European Economic Review*, 40, 603-615

[19] Drehmann, M., Oechssler, J. and Roider, A. (2005), "Herding and contrarian behavior in financial markets: an internet experiment", *The American Economic Review*, 95 (5), 1403-1426

[20] Dufour, A. and Engle, R. (2000), "Time and the price impact of a trade", *The Journal of Finance*, 55 (6), 2467-2498

[21] Easley, D., Kiefer, N., O'Hara, M., and Paperman, J. (1996), "Liquidity, information, and infrequently traded stocks", *The Journal of Finance*, 51 (4), 1405-1436

[22] Eyster, E. and Rabin, M. (2005), "Cursed equilibrium", *Econometrica*, 73 (5), 1623-1672

[23] Frazzini, A. and Lamont, O. (2006), "The earnings announcement premium and trading volume", mimeo, NBER

[24] Hasbrouck, J. (1991), "Measuring the information content of stock trades", *The Journal of Finance*, XLVI (1), 179-207

[25] Hong, H. and Stein, J. (1999), "A unified theory of underreaction, momentum trading, and overreaction in asset markets", *The Journal of Finance* 54, 2143–2184

[26] Hirshleifer, D. and Teoh, S.H. (2003), " Herd behavior and cascading in capital markets: a review and synthesis", *European Financial Management*, 9 (1), 25-66

[27] Ivanov, A., Levin, D., and Peck, J. (2009), "Hindsight, foresight, and insight: an experimental study of a small-market investment game with common and private values", *American Economic Review*, 99, 1484-1507

[28] Ivanov, A., Levin, D., and Peck, J. (2013), "Behavioral biases in endogenous-timing herding games: an experimental study", *Journal of Economic Behavior and Organization,* 87, 25-34

[29] Kandel, E. and Pearson, N. (1995), "Differential interpretation of public signals and trade in speculative markets", *Journal of Political Economy*, 103 (4), 831-872

[30] Kahneman, D. and Tversky, A. (1979), "Prospect theory: an analysis of decision under risk", *Econometrica*, 47 (2), 263-292

[31] Kahneman, D. and Tversky, A. (1992), "Advances in prospect theory: cumulative representation of uncertainty", *Journal of Risk and Uncertainty*, 5, 297-323

[32] Kirchler, M. (2009), "Underreaction to fundamental information and asymmetry in mispricing between bullish and bearish markets. An experimental study", *Journal of Economic Dynamics & Control*, 33, 491-506

[33] Kendall, C. (2015), "The time cost of information in financial markets", mimeo, University of Southern California

[34] Koszegi, B. and Rabin, M. (2006), "A model of reference-dependent preferences", *The Quarterly Journal of Economics*, CXXI (4), 1133-1165

[35] Lee, I. (1998), "Market crashes and information avalanches", *Review of Economic Studies*, 65, 741-759

[36] McKelvey, R. and Palfrey, T (1995), "Quantal response equilibrium for normal form games", *Games and Economic Behavior*, 10, 6-38

[37] McKelvey, R. and Palfrey, T (1998), "Quantal response equilibrium for extensive form games", *Experimental Economics*, 1, 9-41

[38] Nagel, R. (1995), "Unraveling in guessing games: an experimental study", *American Economic Review*, 85 (5), 1313-1326

[39] Park, A. and Sgroi, D. (2012), "Herding, contrarianism, and delay in financial market trading", *European Economic Review*, 56, 1020-1037

[40] Pedersen, L. (2009), "When everyone runs for the exit", *The International Journal of Central Banking*, 5, 177-199

[41] Pettit, J., Hewitt, J. and Oprea, R. (2013), "Redwood: software for graphical, browser-based experiments", mimeo, University of California Santa Barbara

[42] Romer, D. (1993), "Rational asset-price movements without news", *American Economic Review*, 83 (5), 1112-1130

[43] Smith, L. (1997), "Do rational traders frenzy?", mimeo, Massachusetts Institute of Technology

[44] Smith, V, Suchanek, G., and Williams, A. (1988), "Bubbles, crashes and endogenous expectations in experimental spot asset markets", *Econometrica*, 56 (5), 1191-1151

[45] Sgroi, D. (2003), "The right choice at the right time: a herding experiment in endogenous time", *Experimental Economics*, 6, 159-180

[46] Shachat, J. and Srinivasan, A. (2011), "Informational price cascades and non-aggregation of asymmetric information in experimental asset markets", mimeo , SSRN

[47] Stahl, D. and Wilson, P. (1994), "Experimental evidence on players' models of other players", *Journal of Economic Behavior and Organization*, 35, 309-327

[48] Ziegelmeyer, A., My, K.B., Vergnaud, J.C., and Willinger, M. (2005), "Strategic delay and rational imitation in the laboratory", mimeo, Max Planck Institute of Economics

Figures

Figure 1: Cumulative Distribution of Trades Occurring at $t \leq t'$ in Main Treatments

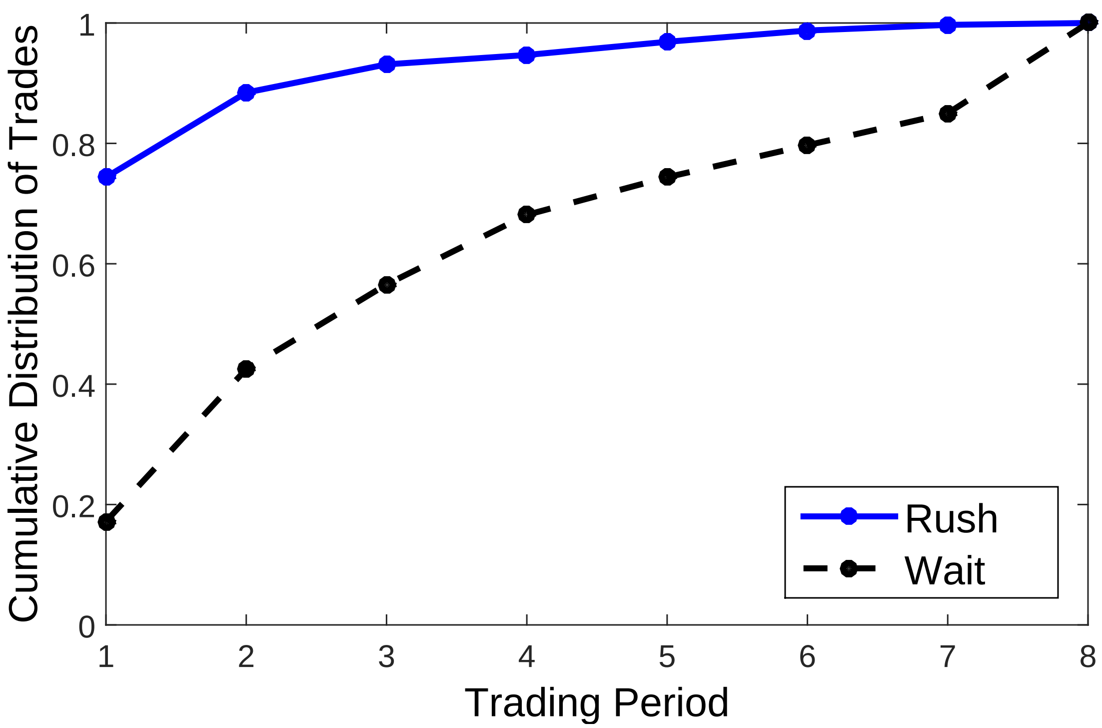

Figure 2: Trade Timing in the Diagnostic Model

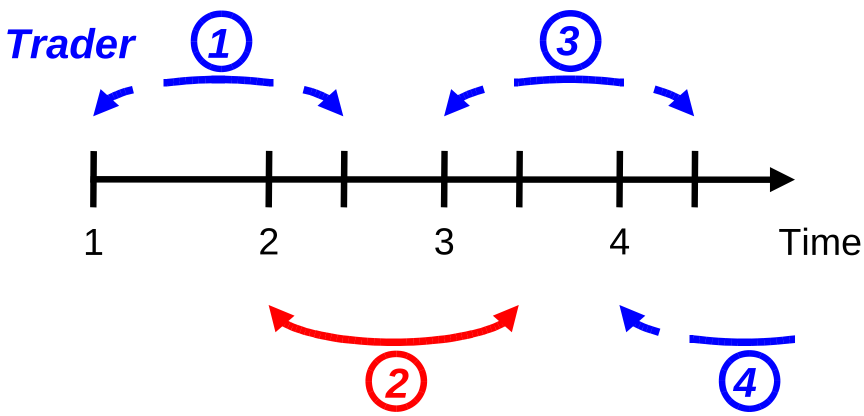

Note: Arrows reflect the two periods in which each trader may trade. The second trader is highlighted for exposition.

Figure 3: Proportion of Rushed Trades In Diagnostic Treatments

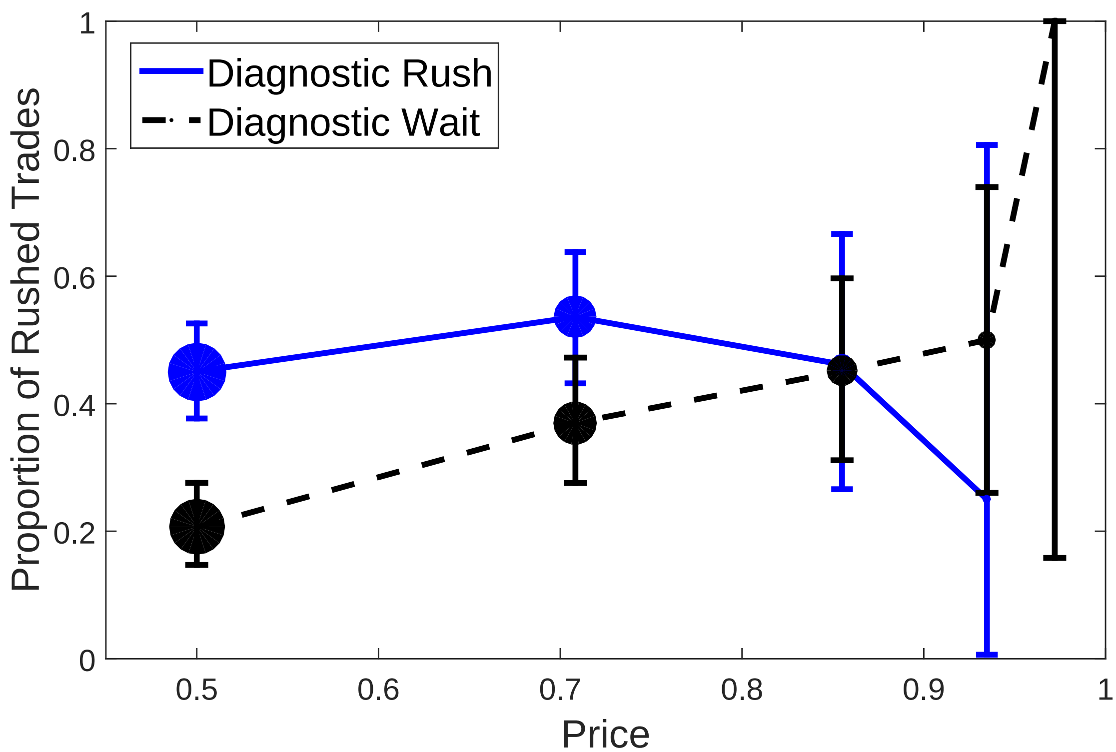

Tables

Table 1: Main Model Treatments

Treatment Name	$\underline{q}$	$\overline{q}$	q_P	Subjects	Periods
Rush (R)	$\frac{3}{4}$	1	$\frac{17}{24}$	$n = 8$	$T = 8$
Wait (W)	$\frac{13}{24}$	1	$\frac{17}{24}$	$n = 8$	$T = 8$

Table 2: Main Model Trading Results

Treatment	Rational	Herding	Contrarian	Irrational
R	93.8% (300)	3.4% (11)	1.9% (6)	0.9% (3)
W	77.2% (247)	15.9% (51)	5.9% (19)	0.9% (3)

Notes: Results reported for last 10 trials. Number of observations in parentheses: 320 total observations per treatment.

Table 3: Determinants of Rushed Trades in the Wait Treatment

	Wait	Wait (τ-herding types)
PubDiff1	0.14***	0.35***
	[0.05]	[0.11]
PubDiff2	0.12*	0.30
	[0.07]	[0.19]
PubDiff3	0.02	0.24
	[0.07]	[0.23]
PubDiff4	-0.08*	-0.03
	[0.04]	[0.10]
PubDiff5	-0.01	0.14
	[0.12]	[0.35]
PubDiff6	-0.06	0.05
	[0.09]	[0.20]
Price (p')	0.02	0.01
	[0.15]	[0.37]
Extreme Signal	0.03	0.10**
	[0.03]	[0.04]
Period	0.08***	0.08***
	[0.01]	[0.03]

Notes: Dependent variable is a dummy variable: 1 indicates a rushed trade. 1077 observations (361 when restricted to τ-herding types). Logit marginal effects reported. Subject and trial fixed effects are included. Robust standard errors in brackets. Significance at the 10% level is represented by *, at the 5% level by **, and at the 1% level by ***.

Table 4: Frequency of Subjects of Each Type In the Wait Treatment

Type	Percent of Subjects (multiple)	(prioritized)
% Rational	0.0-6.3%	**6.3%**
% Simplistic	21.9-37.5%	**37.5%**
% τ-Herding	34.4-56.3%	**34.4%**
% τ-Contrarian	6.3-12.5%	**6.3%**
% Unclassified	15.6%	**15.6%**
% Exact	53.1%	
% Ambiguous	21.9%	

Notes: 32 subjects total. Subjects may match multiple types because I allow for one mistake. Results of assigning a single type through a prioritization scheme are shown in bold.

Table 5: Correlation of Trading Returns in the Wait Treatment

Trading Period	1	2	3	4	5	6	7
Spearman Correlation Coefficient	0.13	0.73***	0.05	0.27	-0.13	-0.03	0.15
	[0.42]	[0.00]	[0.74]	[0.09]	[0.44]	[0.87]	[0.35]

Notes: Correlation is with respect to return due to first public signal. p-value of two-tailed t-test included in brackets. Each correlation is over 40 trials (last 10 trials of each of 4 sessions).

Table 6: Private Signals and Rushed Trades in the Wait Treatment

	Wait
PubDiff1	-0.00
	[0.05]
PubDiff2	0.10
	[0.08]
PubDiff3	-0.01
	[0.08]
PubDiff4	-0.10**
	[0.04]
PubDiff1 * Extreme Signal	0.26**
	[0.10]
PubDiff2 * Extreme Signal	-0.01
	[0.07]
PubDiff3 * Extreme Signal	0.07
	[0.14]
PubDiff4 * Extreme Signal	0.12
	[0.20]
Extreme Signal	-0.05
	[0.05]

Notes: Dependent variable is a dummy variable: 1 indicates a rushed trade. 1077 observations. Logit marginal effects reported. Subject and trial fixed effects are included. Robust standard errors in brackets. Significance at the 10% level is represented by *, at the 5% level by **, and at the 1% level by ***.

Table 7: Diagnostic Model Treatments

Treatment Name	$\underline{q}$	$\overline{q}$	Subjects	Periods
Diagnostic Rush (DR)	$\frac{17}{24}$	$\frac{16}{24}$	$n = 6$	$T = 6$
Diagnostic Wait (DW)	$\frac{13}{24}$	$\frac{17}{24}$	$n = 6$	$T = 6$

Table 8: Diagnostic Treatment Trading Results

Treatment	Rational	Herding	Contrarian	Irrational
Diagnostic Rush (DR)	84.5% (284)	5.7% (19)	6.5% (22)	3.3% (11)
Diagnostic Wait (DW)	83.9% (282)	6.3% (21)	6.9% (23)	3.0% (10)

Notes: Results reported for last 14 trials. Number of observations in parentheses: 336 total observations per treatment.

Table 9: Diagnostic Model Timing Results

Treatment	Frequency of Rush
Diagnostic Rush (DR)	47.6% (147)
Diagnostic Wait (DW)	31.3% (105)

Notes: Results reported for last 14 trials. Number of observations in parentheses: 309 (336) total observations in DR (DW). Rushing corresponds to rational behavior in DR, but not in DW.

Table 10: Determinants of Rushed Trades in the Diagnostic Treatments

	Diagnostic Rush	Diagnostic Wait	Diagnostic Rush (τ-herding types)	Diagnostic Wait (τ-herding types)
Private Belief	0.30	1.69***	1.13**	1.71***
	[0.27]	[0.25]	[0.47]	[0.35]
Previous Rushed	-0.12	-0.04	-0.04	-0.10
	[0.07]	[0.07]	[0.07]	[0.06]

Notes: Dependent variable is a dummy variable: 1 indicates a rushed trade. 308 (294) observations in DW (DR). 132 (84) observations when restricted to τ-herding types. Logit marginal effects reported. Subject and trial fixed effects are included. Robust standard errors in brackets. Significance at the 10% level is represented by *, at the 5% level by **, and at the 1% level by ***.

Table 11: Frequency of Subjects of Each Type in the Diagnostic Treatments

Type	Diagnostic Rush (multiple)	Diagnostic Rush (prioritized)	Diagnostic Wait (multiple)	Diagnostic Wait (prioritized)
% Rational	8.3%	**8.3%**	0-16.7 %	**16.7%**
% Simplistic	0.0%	**0.0%**	4.2%	**4.2%**
% τ-Herding	29.1%	**29.1%**	50-66.7%	**50.0%**
% τ-Contrarian	8.3%	**8.3%**	8.3-33.3%	**8.3%**
% Unclassified	54.1%	**54.1%**	20.8%	**20.8%**
% Exact	29.1%		54.1%	
% Ambiguous	0.0%		25.0%	

Notes: 32 subjects for each of DR and DW. Subjects may match multiple types because I allow for one mistake. Results of assigning a single type through the prioritization scheme are shown in bold.

Appendix

A. Learning

Figure A1 plots the evolution of rational timing decisions in each treatment. The clearest evidence of learning is observed in the R treatment, where robust convergence to a rational panic is observed. In the DW treatment, learning is also evident: subjects learn that an additional signal is valuable. On the other hand, they do not appear to learn this same lesson in the W treatment. It may be that the large number of opportunities to trade makes learning more difficult: if one never experiments with waiting until the final period, one doesn't observe its benefit. Finally, in the DR treatment, learning is particularly difficult because waiting is only costly if either the previous trader waited or the subsequent trader rushes and the trade(s) move prices adversely. If one's successors mostly wait, it may take time to learn that prices on average move adversely. Learning may have perhaps been easier had subjects maintained the same position in the sequence in treatments DR and DW, but the concern in this case is that subjects would be learning about a particular subject's behavior rather than about their environments.

B. Non-equilibrium Informational Losses

Non-equilibrium behavior exacerbates the informational losses predicted by rational theory in all four treatments. I measure informational losses by comparing final prices in each trial to actual asset values. Figure B1 plots the empirical cumulative distribution function (cdf) of the absolute value of the difference in these two values for each of the four treatments. In each case, the difference between the asset value and the theoretical final price, had all traders acted rationally, is compared to the difference between the asset value and the actual final price in the data. In the two treatments in which rushing is optimal, R and DR, we see

Figure A1: Proportion of Rational Timing Decisions as Trials Progress

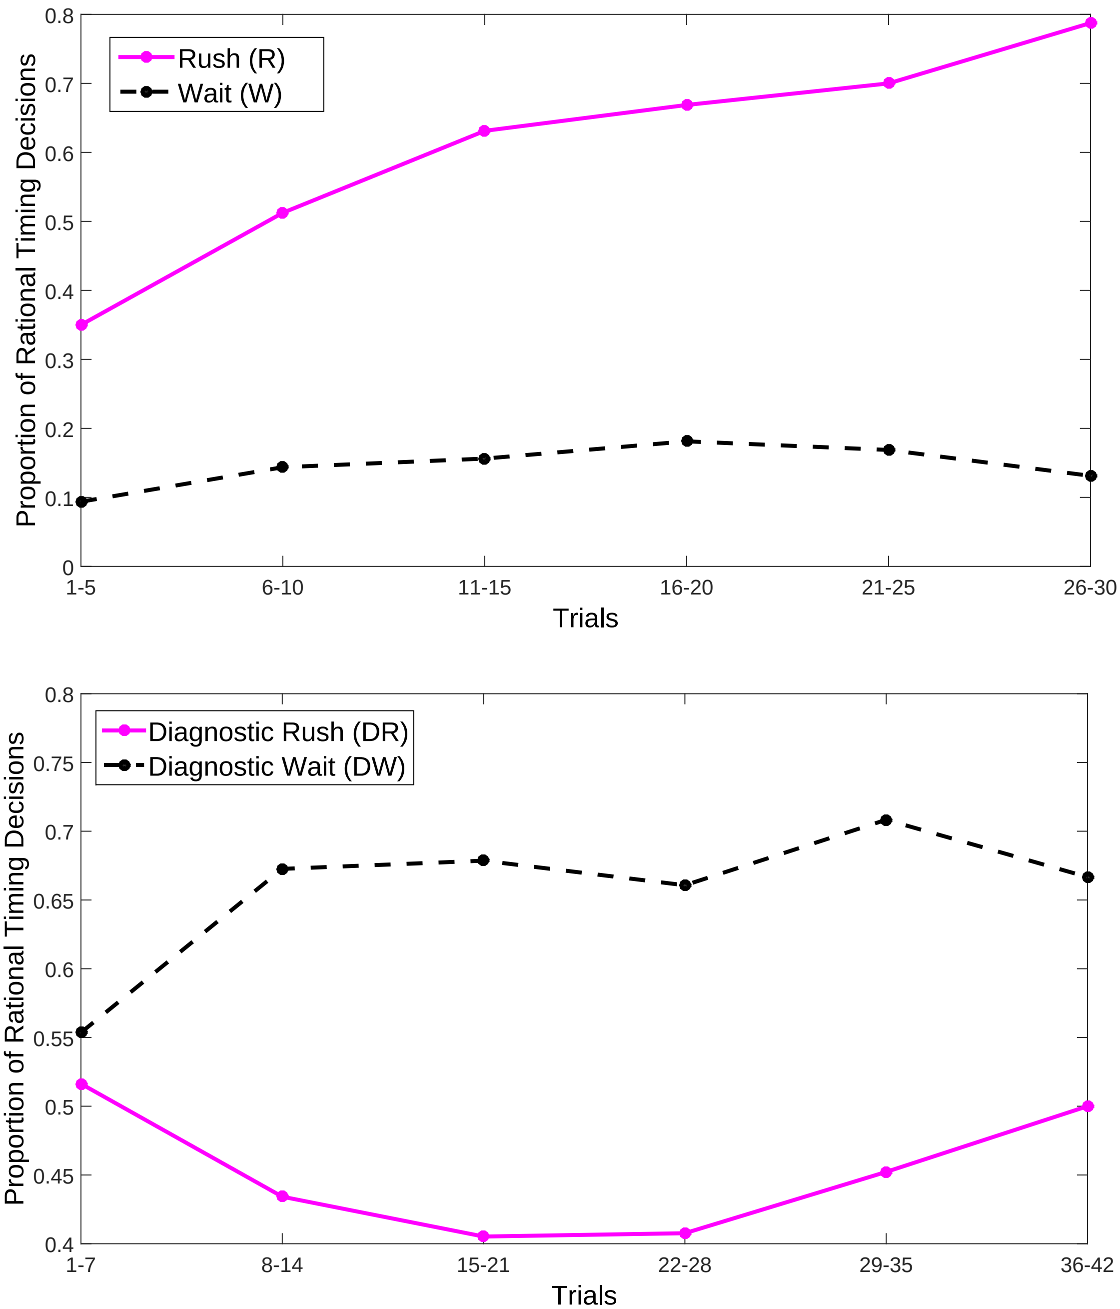

Figure B1: CDFs of Difference Between Final Prices and Asset Values by Treatment

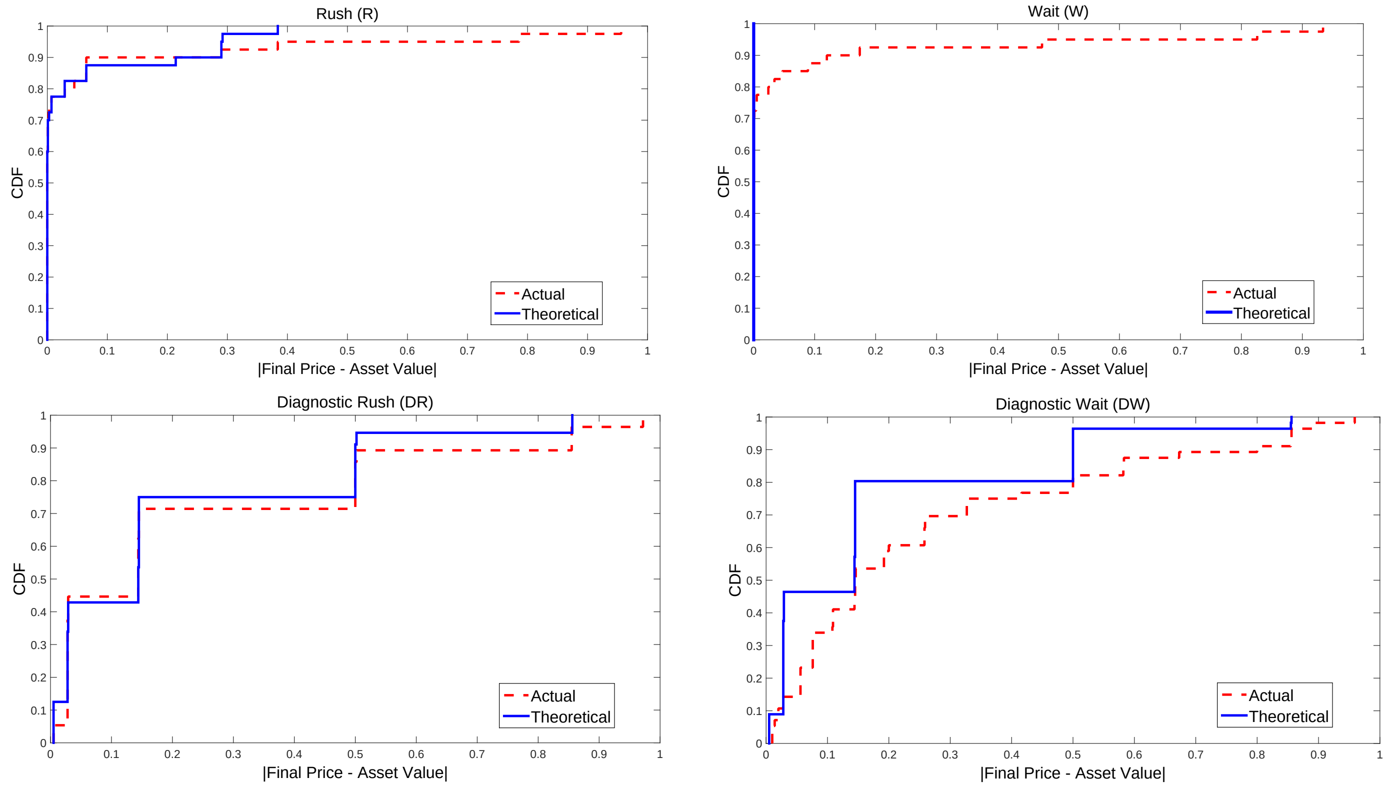

that the theoretical and actual losses are very similar: differences are mostly due to trades not revealing private information. Therefore, when rushing is rational, almost all information losses are due to rational behavior.[48] Importantly, non-equilibrium waiting does not cause additional information to be aggregated because, in DR, waiting to obtain an additional signal does not reveal additional information and in R, only one subject in one trial waited long enough to obtain this information.

In the treatments in which waiting is optimal, we observe large informational losses due to non-equilibrium behavior. In treatment W, where all traders should obtain perfect information, informational losses are particularly surprising because, should any *one* of the eight traders wait to obtain perfect information, the final price would reflect the true asset value. Thus, in the trials in which information is lost, *no* trader waited. In treatment DW, where traders determine all information aggregation, we observe the largest losses due to non-equilibrium behavior. On average, the observed final price differs from the theoretical by 9.6%. Average differences in the other three treatments are: R (2.6%), W (6.8%), and DR (3.5%).

[48]In treatment R, final prices are generally very close to the true asset value even though traders forgo perfect information, because final prices also reflect the information contained in the public signals.

C. Estimation of Individual Thresholds

The τ-herding heuristic posits an individual belief threshold, τ_i. If behavior is perfectly consistent with the heuristic, then we expect to observe waiting at all beliefs between $1 - \tau_i$ and τ_i, and rushing for beliefs outside this range. Focusing on a belief, $b_{i,t} > \frac{1}{2}$, the probability of observing a rush decision is given by

$$Pr(R) = Pr(b_{i,t} + \varepsilon_{i,t} > \tau_i)$$

$$\iff Pr(R) = Pr(\varepsilon_{i,t} > \tau_i - b_{i,t}) \tag{1}$$

where $\varepsilon_{i,t}$ captures any observed component that contributes to the decision to rush. Equation (1) assumes errors are small so that if a rush decision is observed for $b_{i,t} > \frac{1}{2}$, it is due to the upper threshold, τ_i, being surpassed (as opposed to the lower threshold, $1 - \tau_i$). It suggests a simple logit model:

$$Pr(R) = \gamma_i(\tau_i - b_{i,t}) + \varepsilon_{i,t}$$

where the heuristic prediction is that $\gamma_i \to -\infty$ and $\tau_i < 1$. $\gamma_i \to -\infty$ occurs if an individual uses a precise threshold but, given the discrete nature of possible beliefs in the data, the threshold can not be pinned down exactly. We therefore expect only for γ_i to be large and negative. Assuming symmetric behavior around $b_{i,t} = \frac{1}{2}$, I obtain the individual τ_i estimates using

$$Pr(R) = \gamma_i(\tau_i - \max\{b_{i,t}, 1 - b_{i,t}\}) + \varepsilon_{i,t}$$

where the regression is run for each individual using observations from the last 10 (14 in the Diagnostic treatments) trials. Using the entire set of trials produces similar results in terms of the number of individuals with thresholds significantly less than one.

For Online Publication

Online Appendix

A. Analysis Details and Omitted Proofs

A1. Preliminaries (Applicable to Both Models)

The overall strategy to prove Propositions 1, 2, D1, and D2 is to prove a lemma which allows the results of Kendall (2015) to be applied. The main difference between both the main model and the diagnostic model from the model in Kendall (2015) is that traders here have private information when they make their timing decision. As such, it is possible that traders with different initial private signals (which I refer to as their type) make different timing decisions, revealing information to the market.

The results in this section apply to both models. I use the notation of the main model but the results are easily translated. The solution concept is sequential equilibrium and I restrict attention to Markov strategies which are a function only of the payoff-relevant state.[49] To determine the optimal timing decision(s) of a trader, one must determine the benefit from not trading in the current period, but instead trading in the future. This benefit in turn depends upon the optimal trading strategies that the trader follows conditional on reaching each particular information set in the future. I denote the equilibrium probability that a trader of type $\underline{s} = x$ for $x \in \{0,1\}$ does not trade in the current period as β_x. For given values of β_0 and β_1, Lemma A1 provides the optimal trading strategies in the current period and the sequentially rational trading strategies at any future information set on the equilibrium path.[50,51] $\hat{q}$ is used to denote the quality of information received in the future period. When no new private information is received, it can be modeled as a hypothetic signal, $\overline{s}_i$, with quality $\frac{1}{2}$. In the Diagnostic model, $\hat{q} = \overline{q}$ always. In the main model, $\hat{q} = \overline{q}$ in period T but $\hat{q} = \frac{1}{2}$ otherwise.

Lemma A1: *In any equilibrium:*

1. *Any trader who trades in the current period buys if $\underline{s}_i = 1$ and sells if $\underline{s}_i = 0$.*
2. *Any trader who waits until a future period and then trades, trades as follows::*

   (a) *buys if $\underline{s}_i = 1$ and $\overline{s}_i = 1$; sells if $\underline{s}_i = 0$ and $\overline{s}_i = 0$*
   (b) *buys (sells) if $\underline{s}_i = 0$, $\overline{s}_i = 1$, and $g_0(\underline{q}, \hat{q}) \equiv (1-\underline{q})\hat{q}NT_0 - \underline{q}(1-\hat{q})NT_1 \geq (<) 0$*
   (c) *buys (sells) if $\underline{s}_i = 1$, $\overline{s}_i = 0$, and $g_1(\underline{q}, \hat{q}) \equiv \underline{q}(1-\hat{q})NT_0 - (1-\underline{q})\hat{q}NT_1 > (\leq) 0$*

   where $NT_0 = (1-\underline{q})\beta_1 + \underline{q}\beta_0$, $NT_1 = \underline{q}\beta_1 + (1-\underline{q})\beta_0$ are shorthand for the probabilities, from the market maker's persepective, of observing no trade in the current period conditional on $V = 0$ and $V = 1$, respectively.

Proof of Lemma A1:

The proof when rushing is similar to the proof when waiting, only simpler, so is omitted for brevity. Let $\hat{a}$ denote any information that becomes public other than from trader i herself not trading (denoted NT), and abbreviate $Pr(\hat{a}|V = y)$ as $\hat{a}_y$, for $y \in \{0,1\}$. A trader who waits buys if her expected value of the asset conditional on her information in the future period is greater than the future price. Denoting the current history, H_t, the

[49] Because the market maker is not a strategic player, assumptions on off-equilbrium price formation must also be imposed. I assume prices do not change after an off-equilibrium timing decision and are updated in a manner consistent with sequential equilibrium after an off-equilibrium trade. See Kendall (2015) for a more detailed discussion and a formal equilibrium definition in a related model.

[50] I address the issue of off-equilibrium beliefs and strategies below.

[51] I implicitly assume that the timing strategies of the two types of traders differ only in the current period such that, when β_0 and β_1 differ, the same information is revealed to the market on every future path. I show below that one need not consider the more general case.

current price p_t, and the future price, $Pr(V=1|\hat{a}, NT, H_t)$, a trader buys if

$$
\begin{aligned}
& Pr(V=1|\underline{s}_i, \overline{s}_i, \hat{a}, H_t) > Pr(V=1|\hat{a}, NT, H_t) \\
\iff & \frac{p_t Pr(\underline{s}_i|V=1) Pr(\overline{s}_i|V=1)\hat{a}_1}{p_t Pr(\underline{s}_i|V=1) Pr(\overline{s}_i|V=1)\hat{a}_1 + (1-p_t) Pr(\underline{s}_i|V=0) Pr(\overline{s}_i|V=0)\hat{a}_0} \\
& > \frac{p_t \hat{a}_1 Pr(NT|V=1)}{p_t \hat{a}_1 Pr(NT|V=1) + (1-p_t)\hat{a}_0 Pr(NT|V=0)} \\
\iff & Pr(\underline{s}_i|V=1) Pr(\overline{s}_i|V=1) NT_0 > Pr(\underline{s}_i|V=0) Pr(\overline{s}_i|V=0) NT_1 \qquad (2)
\end{aligned}
$$

where the first equivalence follows from applying Bayes' rule to each side of the inequality and using the fact that the public belief is the initial price, $Pr(V=1|H_t) = p_t$. Using (2), a trader with $\underline{s}_i = 1$ and $\overline{s}_i = 1$ buys if

$$
\begin{aligned}
& \underline{q}\hat{q} NT_0 > (1-\underline{q})(1-\hat{q}) NT_1 \\
\iff & \underline{q}(1-\underline{q})(2\hat{q}-1)\beta_1 + \left(\underline{q}^2\hat{q} - (1-\underline{q})^2(1-\hat{q})\right)\beta_1 > 0
\end{aligned}
$$

which is true for all parameterizations. Similarly, a trader with $\underline{s}_i = 0$ and $\overline{s}_i = 0$ always sells. Finally, the conditions in Lemma A1 under which traders with $\underline{s}_i = 1$ and $\overline{s}_i = 0$, or $\underline{s}_i = 0$ and $\overline{s}_i = 1$, buy or sell are easily obtained by substituting for the appropriate probabilities into (2).□

Given Lemma A1, I derive a general function for the benefit from waiting (the expected profit from future trades less the profit if one trades at $t = 1$), assuming that the future information sets in which trades occur are on the equilibrium path. When this benefit is positive (negative), a trader optimally waits (rushes). The benefit depends upon the probability each type of trader waits because information is revealed if these probabilities differ across types. In equilibrium, the two probabilities must be consistent with the sign of the benefit of each type of trader. I denote the benefit from waiting, $B_x(p_t, \beta_0, \beta_1)$, for traders with $\underline{s}_i = x$.

The benefit from waiting also depends upon any public information revealed by others' strategies and any public signals. I use generic notation to denote the possible informative events because several properties of the benefit can be established for any informative event. I denote a generic event that occurs between the current and future periods, $\hat{a}$, and the set of possible such events as A, so that $\hat{a} \in A$. I also abbreviate $Pr(\hat{a}|V=y)$ as $\hat{a}_y$, for $y \in \{0,1\}$, so that, when summing over all possible events, $\sum_{\hat{a}\in A}\hat{a}_0 = \sum_{\hat{a}\in A}\hat{a}_1 = 1$.[52]

Consider first the expected profit of a trader who trades in the current period, t, with $\underline{s}_i = 1$. From Lemma A1, she buys so that her profit is given by

$$
Pr(V=1|\underline{s}_t = 1, \underline{H}_t) - p_t \iff \frac{p_t \underline{q}}{p_t\underline{q} + (1-p_t)(1-\underline{q})} - p_t \iff \frac{p_t(1-p_t)(2\underline{q}-1)}{p_t\underline{q} + (1-p_t)(1-\underline{q})}
$$

The profit for a trader with $\underline{s}_i = 0$ is calculated similarly, using the fact that she sells. The current profit for both types of traders can be written $\frac{p_t(1-p_t)(2\underline{q}-1)}{Pr(\underline{s}_i)}$.

[52] As an example, an event could be the revelation of a public signal at some time, r, in which case $A = \{s_{P,r} = 1, s_{P,r} = 0\}$ is the set of possible events. Then, $\sum_{\hat{a}\in A}\hat{a}_1 = \sum_{\hat{a}\in A}\hat{a}_0 = q_P + 1 - q_P = 1$.

I calculate the expected profit from future trades for a trader with $\underline{s}_i = 1$ who buys when receiving $\overline{s}_i = 1$ and sells when $\overline{s}_i = 0$ ($g_1(\underline{q}, \hat{q}) \leq 0$). The other cases are calculated similarly and are omitted for brevity. The expected profit is

$$\sum_{\hat{a} \in A} \{Pr(\overline{s}_i = 1 \& \hat{a} | \underline{s}_i = 1) \left(Pr(V = 1 | \underline{s}_i = 1, \overline{s}_i = 1, \hat{a}) - Pr(V = 1 | NT, \hat{a})\right) \\ + Pr(\overline{s}_i = 0 \& \hat{a} | \underline{s}_i = 1) \left(Pr(V = 1 | NT, \hat{a}) - Pr(V = 1 | \underline{s}_i = 1, \overline{s}_i = 0, \hat{a})\right)\}$$

where all probabilities are also conditional on H_t. Here, I sum over each possible generic combination of events that result from the public events. The first term corresponds to the profit from buying the asset after receiving $\overline{s}_i = 1$ and the second term from selling after receiving $\overline{s}_i = 0$. Using Bayes' rule and the independence of signals, the above becomes

$$\sum_{\hat{a} \in A} \left\{ \frac{Pr(\overline{s}_i = 1 \& \underline{s}_i = 1 \& \hat{a})}{Pr(\underline{s}_i = 1)} \left(\frac{p_t \underline{q} \hat{q} Pr(\hat{a} | V = 1)}{Pr(\overline{s}_i = 1 \& \underline{s}_i = 1 \& \hat{a})} - \frac{p_t NT_1 Pr(\hat{a} | V = 1)}{Pr(NT \& \hat{a})} \right) \right. \\ \left. + \frac{Pr(\overline{s}_i = 0 \& \underline{s}_i = 1 \& \hat{a})}{Pr(\underline{s}_i = 1)} \left(\frac{p_t NT_1 Pr(\hat{a} | V = 1)}{Pr(NT \& \hat{a})} - \frac{p_t \underline{q}(1 - \hat{q}) Pr(\hat{a} | V = 1)}{Pr(\overline{s}_i = 0 \& \underline{s}_i = 1 \& \hat{a})} \right) \right\}$$

with NT_0 and NT_1 as in Lemma A1. Combining each of the terms in the first and second expressions, canceling and factoring out common terms gives

$$\frac{p_t(1 - p_t)}{Pr(\underline{s}_i = 1)} \sum_{\hat{a} \in A} \frac{\hat{a}_0 \hat{a}_1}{Pr(NT \& \hat{a})} \left(\underline{q} \hat{q} NT_0 - NT_1(1 - \hat{q})(1 - \underline{q}) + NT_1 \hat{q}(1 - \underline{q}) - \underline{q}(1 - \hat{q}) NT_0 \right) \\ \iff \frac{p_t(1 - p_t)}{Pr(\underline{s}_i = 1)} \sum_{\hat{a} \in A} \frac{\hat{a}_0 \hat{a}_1}{Pr(NT \& \hat{a})} \left((2\hat{q} - 1)(\underline{q} NT_0 + (1 - \underline{q}) NT_1) \right)$$

Finally, subtracting the profit at t from the expected profit and factoring out common terms gives the benefit.

Considering both types of traders and the possible future trading strategies in Lemma A1, we have

$$B_x(p_t, \beta_0, \beta_1) = \frac{p_t(1 - p_t)}{Pr(\underline{s}_i = x)} \left[\sum_{\hat{a} \in A} \frac{\hat{a}_0 \hat{a}_1 f(\underline{q}, \hat{q}, \beta_0, \beta_1)}{Pr(\hat{a} \& NT)} - (2\underline{q} - 1) \right] \tag{3}$$

where

$$f(\underline{q}, \hat{q}, \beta_0, \beta_1) = \begin{cases} (2\hat{q} - 1)(\underline{q} NT_0 + (1 - \underline{q}) NT_1) & if\ \underline{s}_i = 1 \ \&\ g_1(\underline{q}, \hat{q}) \leq 0 \\ \underline{q} NT_0 - (1 - \underline{q}) NT_1 & if\ \underline{s}_i = 1 \ \&\ g_1(\underline{q}, \hat{q}) > 0 \\ \underline{q} NT_1 - (1 - \underline{q}) NT_0 & if\ \underline{s}_i = 0 \ \&\ g_0(\underline{q}, \hat{q}) < 0 \\ (2\hat{q} - 1)(\underline{q} NT_1 + (1 - \underline{q}) NT_0) & if\ \underline{s}_i = 0 \ \&\ g_0(\underline{q}, \hat{q}) \geq 0 \end{cases}$$

Note that $\hat{q}$ in (3) may depend on the information set reached (and therefore on the event, $\hat{a}$, that occurs). Given the generic formula for the benefit, Lemma A2 establishes that both types of traders must follow the same timing strategies in any equilibrium. Importantly, this result holds without any assumptions on off-equilibrium beliefs or strategies.

Lemma A2: *In any equilibrium, traders with* $\underline{s}_t = 0$ *and* $\underline{s}_t = 1$ *must follow the same timing strategy at every history.*

Proof of Lemma A2:

I provide the proof for the main model. The proof in the Diagnostic model uses the same reasoning but is simpler because traders wait at most one period.

The proof uses backwards induction. Beginning in period $T-1$, I show that the claim holds. Then, I show that it holds for an arbitrary period, t, given that it holds for all subsequent periods. Within a period, the proof is by contradiction. Assume that there exists an equilibrium in which traders wait with different probabilities ($\beta_1 \neq \beta_0$) in period t after some history, H_t. Then, in all future periods, the two types of traders arrive with these same probabilities, using the fact that in all periods greater than t they use the same timing strategy.[53]

Without loss of generality, assume $\beta_1 > \beta_0$. Given $\beta_1 > \beta_0$, we must have $\beta_1 \in (0,1]$ and therefore $B_1(p_t, \beta_0, \beta_1) \geq 0$ (type $1's$ benefit must be weakly positive in equilibrium if she waits with positive probability). I establish below that $\beta_1 > \beta_0$ and $B_1(p_t, \beta_0, \beta_1) \geq 0$ together imply $B_0(p_t, \beta_0, \beta_1) > 0$. However, $\beta_1 > \beta_0$ also implies $\beta_0 \in [0,1)$ and therefore $B_0(p_t, \beta_0, \beta_1) \leq 0$ (if type $0's$ benefit is strictly positive, she waits with probability one), a contradiction. Therefore, we must have $\beta_1 = \beta_0$ in any equilibrium.

I now establish $\beta_1 > \beta_0$ and $B_1(p_t, \beta_0, \beta_1) \geq 0$ together imply $B_0(p_t, \beta_0, \beta_1) > 0$. One can see from the general form of the benefit in (3), that the sign of $B_x(p_t, \beta_0, \beta_1)$ is determined by the term in square brackets and that the only difference between the bracketed term across types is due to differences in the functions, $f(\underline{q}, \hat{q}, \beta_0, \beta_1)$, that depend on the event that occurred ($\hat{q}$ can differ across events). I show that for any event, $f(\underline{q}, \hat{q}, \beta_0, \beta_1)$ is strictly greater for type 0 than for type 1 whenever $\beta_1 > \beta_0$, so that if $B_1(p_t, \beta_0, \beta_1) \geq 0$, we must have $B_0(p_t, \beta_0, \beta_1) > 0$.

First consider an information set in which no new information is received ($\hat{q} = \frac{1}{2}$). In this case, $g_0(\underline{q}, \hat{q}) < 0$ and $g_1(\underline{q}, \hat{q}) > 0$. Comparing $f(\underline{q}, \hat{q}, \beta_0, \beta_1)$ in $B_0(p_t, \beta_0, \beta_1)$ relative to $B_1(p_t, \beta_0, \beta_1)$ gives $\underline{q}NT_1 - (1-\underline{q})NT_0 > \underline{q}NT_0 - (1-\underline{q})NT_1 \iff NT_1 - NT_0 > 0 \iff (\beta_1 - \beta_0)(2\underline{q} - 1) > 0$ so $f(\underline{q}, \overline{q}, \beta_0, \beta_1)$ is strictly greater in $B_0(p_t, \beta_0, \beta_1)$.

Now consider an information set in which new information is received ($\hat{q} = \overline{q}$). There are four possible cases depending upon the optimal strategies of buying and selling from Lemma A1.

1. $g_1(\underline{q}, \overline{q}) \leq 0$ and $g_0(\underline{q}, \overline{q}) \geq 0$. The comparison of $f(\underline{q}, \overline{q}, \beta_0, \beta_1)$ in $B_0(p_t, \beta_0, \beta_1)$ relative to $B_1(p_t, \beta_0, \beta_1)$ is $\underline{q}NT_1 + (1-\underline{q})NT_0 > \underline{q}NT_0 + (1-\underline{q})NT_1 \iff (NT_1 - NT_0)(2\underline{q}-1) > 0 \iff (\beta_1 - \beta_0)(2\underline{q} - 1)^2 > 0$ so $f(\underline{q}, \overline{q}, \beta_0, \beta_1)$ is strictly greater in $B_0(p_t, \beta_0, \beta_1)$.
2. $g_1(\underline{q}, \overline{q}) > 0$ and $g_0(\underline{q}, \overline{q}) < 0$. This case is identical to the $\hat{q} = \frac{1}{2}$ case.

[53] At $T-1$, the two types of traders trivially arrive in period T with probabilities β_0 and β_1.

3. $g_1(\underline{q}, \overline{q}) \leq 0$ and $g_0(\underline{q}, \overline{q}) < 0$. Note that $g_0(\underline{q}, \overline{q}) < 0 \implies \underline{q}NT_1 - (1-\underline{q})NT_0 > (2\overline{q}-1)(\underline{q}NT_1 + (1-\underline{q})NT_0)$. This can be seen algebraically or simply by noting that $B_0(p_t, \beta_0, \beta_1)$ must be larger when t follows her optimal trading strategy instead of the optimal strategy for the other case, $g_0(\underline{q}, \overline{q}) \geq 0$. But, then we have, $\underline{q}NT_1 - (1-\underline{q})NT_0 > (2\overline{q}-1)(\underline{q}NT_1 + (1-\underline{q})NT_0) > (2\overline{q}-1)(\underline{q}NT_0 + (1-\underline{q})NT_1)$ where the second inequality was shown in the first case above, so $f(\underline{q}, \overline{q}, \beta_0, \beta_1)$ is strictly greater in $B_0(p_t, \beta_0, \beta_1)$.
4. $g_1(\underline{q}, \overline{q}) > 0$ and $g_0(\underline{q}, \overline{q}) \geq 0$. $g_0(\underline{q}, \overline{q}) \geq 0$ implies $(2\overline{q}-1)(\underline{q}NT_1 + (1-\underline{q})NT_0) > \underline{q}NT_1 - (1-\underline{q})NT_0$ by optimality of the future trading decision and $\underline{q}NT_1 - (1-\underline{q})NT_0 > \underline{q}NT_0 - (1-\underline{q})NT_1$ as shown for the $\hat{q} = \frac{1}{2}$ case. Thus, $f(\underline{q}, \overline{q}, \beta_0, \beta_1)$ is strictly greater in $B_0(p_t, \beta_0, \beta_1)$.

This completes the claim that $\beta_1 > \beta_0$ and $B_1(p_t, \beta_0, \beta_1) \geq 0$ together imply $B_0(p_t, \beta_0, \beta_1) > 0$ and therefore completes the proof.□

Given that, in equilibrium, both types of traders must follow the same timing strategy, the equilibrium trading strategies of Lemma A1 simplify (setting $\beta_0 = \beta_1$ in the formulas), resulting in Propositions 1 and D1.

The benefit function simplifies as well. $f(\underline{q}, \hat{q}, \beta_0, \beta_1)$ simplifies to $\beta(2\hat{q}-1)$ when $\hat{q} > \underline{q}$ and $\beta(2\underline{q}-1)$ when $\hat{q} < \underline{q}$, reflecting the fact that a trader follows her better quality signal when they are contradictory. The denominator simplifies to $Pr(\hat{a} \& \underline{a}_t = NT) = \beta Pr(\hat{a})$ so that β cancels in the numerator and denominator. When $\overline{q} = \underline{q}$, a trader with contradictory signals is indifferent between trading or not, but under the assumption that such a trader follows her second period signal, $f(\underline{q}, \hat{q}, \beta_0, \beta_1) = \beta(2\hat{q}-1)$ and β cancels in this case as well. Denoting $q = max(\underline{q}, \hat{q})$, one can then write the general benefit function,

$$B_x(p_t) = \frac{p_t(1-p_t)}{Pr(\underline{s}_t = x)} \left[\sum_{\hat{a} \in A} \frac{\hat{a}_0 \hat{a}_1 (2q-1)}{Pr(\hat{a})} - (2\underline{q}-1) \right] \tag{4}$$

Thus far, I have not made any assumption about off-equilibrium beliefs. An assumption is necessary to pin down the benefit a trader gets from deviating when she rushes in equilibrium (the off-equilibrium benefit must be negative in this case). Given the result of Lemma A2, a natural assumption is that any unexpected wait decision is equally likely to be from either type of trader (as it must be on the equilibrium path). I impose this assumption in the remainder of the analysis. The optimal trading strategies after a deviation and the benefit from deviation are given by Lemma A1 and (3) except that beliefs about the probabilities that each type of trader waits replace β_1 and β_0. Under the assumption that any deviation is equally likely to be from either type, the benefit simplifies as it does on the equilibrium path, so that it is given by (4). A nice property of this assumption is that the off-equilibrium benefit does not in fact depend on the beliefs (other than that they are equal across types).

To conclude this section, I note the similarity between (4) and the benefit function in Kendall (2015). The only differences are the scale factor outside of the bracketed term and the fact that q replaces $\underline{q}$, allowing for the possibility that one trades at an information set with no new private information. This similarity implies that the zeros of the benefit function (and hence when it is positive or negative) are identical to those of the benefit function in Kendall (2015) so that one can apply the results of that paper. Of specific importance are Proposition 1 and Lemma C1. Proposition 1 establishes that informative events strictly reduce the benefit from waiting at all prices. Lemma C1 establishes various properties of the benefit function, including properties of its zero-crossings.

A2. Main Model

In this section, I prove Propositions 2 (Proposition 1 was proven in Section A1). I focus on $\overline{q} > \underline{q}$ as in the experiment. For $\overline{q} \leq \underline{q}$, it is easily shown that the unique equilibrium is for traders to trade immediately. Markov strategies in this model depend only upon the price and the number of traders that have traded in prior periods, $\hat{n} \in 0, 1, \ldots, n-1$.

For general parameters, $\underline{q}$,$\overline{q}$, and q_P, a multiplicity of equilibria may exist. Intuitively, the general form of the game is a coordination game where, if others wait, one is willing to wait, but if others rush, one rushes. However, the public signals impose a cost of waiting independent of what the other traders do. In Lemma A3, I use this fact to establish sufficient conditions for the equilibrium outcome to be unique. Define the benefit from waiting from $t = 1$ until $t = T$ when all other traders also wait until $t = T$ as $B_x^W(\frac{1}{2})$ where the initial price, $p_1 = \frac{1}{2}$, as in the experiment. Similarly, define the benefit from waiting from $t = 1$ until $t = T$ when all other traders rush $B_x^R(\frac{1}{2})$.

Lemma A3:

1. *If* $B_x^W(\frac{1}{2}) < 0$ *for* $x \in \{0, 1\}$, *in any equilibrium traders trade at* $t = 1$.
2. *If* $B_x^R(\frac{1}{2}) > 0$ *for* $x \in \{0, 1\}$, *then in the unique equilibrium, traders wait at all histories in periods* $t = 1, 2, \ldots, T-1$ *and trade in period* T.

Proof of Lemma A3:

First note that the signs of $B_0^R(\frac{1}{2})$ and $B_1^R(\frac{1}{2})$ (and of $B_0^W(\frac{1}{2})$ and $B_1^W(\frac{1}{2})$) are identical because they differ by only a scale factor. Therefore, take $x = 1$ for concreteness.

Part 1). The proof is by contradiction. Suppose that an equilibrium exists in which a trader waits at $t = 1$, and denote the benefit from waiting, $B_1^{W*}(\frac{1}{2})$, for some plan of future trades denoted $(*)$. To be a best response, we must have $B_1^{W*}(\frac{1}{2}) > 0$. If $(*)$ involves always waiting until T, then $B_1^{W*}(\frac{1}{2}) = B_1^W(\frac{1}{2}) < 0$, an immediate contradiction. So, suppose instead $(*)$ involves trading at some intermediate trading period, $t = 1, 2, \ldots, T-1$ at some information set I^* and price, p_t. The benefit from waiting at I^* must be less than zero to support the plan $(*)$ in equilibrium. At I^*, less than $T-1$ public signals remain, so by Proposition 1 of Kendall (2015), the benefit from waiting at I^* must be greater than $B_1^{W*}(p_t)$ because the benefit strictly increases when less informative events remain. Also, by the properties established in Lemma C1 of Kendall (2015), if $B_1^{W*}(\frac{1}{2}) > 0$ then $B_1^{W*}(p_t) > 0$ for all p_t. Therefore, the benefit from waiting at I^*must be greater than $B_1^{W*}(\frac{1}{2}) > 0$, a contradiction.

Because all traders i must trade at $t = 1$ even when all other traders wait until i, it remains to check that all trading at $t = 1$ is an equilibrium. By Proposition 1 of Kendall (2015), the benefit from waiting at $t = 1$ is further reduced by the trades of other traders at $t = 1$, so it is in fact an equilibrium. Note that part 1) does not claim equilibrium uniqueness because it leaves open the strategy after off-equilibrium histories in which a trader waits until $t > 1$. In particular, it does *not* say that a trader must trade immediately at every (off-equilibrium) history. In fact, by Proposition 2 of Kendall (2015), as prices approach either 0 or 1, the benefit from waiting until T becomes positive so that there are histories at which waiting becomes optimal, conditional on reaching such a history.

Part 2). If $B_1^R(\frac{1}{2}) > 0$, then $B_1^R(p_t) > 0$ for all p_t according to the properties of any benefit function established in Lemma C1 of Kendall (2015). If $B_1^R(p_t) > 0$ at some p_t, then even if all other traders trade between $t = 1$ and $t = T$ and $T - 1$ public signals remain, trader i would wait until T. By Proposition 1 of Kendall (2015), if some of the other traders have already traded ($\widetilde{n} > 1$) and/or if we have reached a trading period where less than $T-1$ public signals remain, then because less informative events remain, i's benefit from waiting would be strictly greater than $B_1^R(p_t)$ and therefore greater than zero. Therefore, no matter what the strategies of the other traders are or what trading period i is in, i waits until T. Because this holds for all $i \in n$ and at every history, the unique equilibrium is for all traders to wait to trade until T.$\square$

Using Lemma A3, we can establish Proposition 2 by checking that one of the sufficient conditions is satisfied for each parameterization.

Proposition 2, part a). In general, the sets of parameters $(\underline{q}, \overline{q}, q_P, p_1, T)$ that satisfy $B_x^W(p_1) < 0$ for $x \in \{0, 1\}$ can be quite wide and, given the complexity of this benefit function, it is difficult to provide a simple characterization. However, one can evaluate $B_x^W(p_1)$ for the parameters in treatment R: $\underline{q} = \frac{3}{4}$, $\overline{q} = 1$, $q^* = \frac{17}{24}$, $p_1 = \frac{1}{2}$, and $T = 8$. Because only public signals affect $B_x^W(p_1)$, we can rewrite it as

$$B_x^W(p_1) = \frac{p_1(1-p_1)}{Pr(\underline{s}_i = x)} \left[\sum_{k=0}^{T-1} \frac{C_0(k)C_1(k)(2\overline{q}-1)}{p_1 C_1(k) + (1-p_1)C_0(k)} - (2\underline{q}-1) \right] \tag{5}$$

where $C_0(k) = \frac{(T-1)!}{k!(T-1-k)!}(1-q_P)^k q_P^{T-1-k}$ and $C_1(k) = \frac{(T-1)!}{k!(T-1-k)!} {q_P}^k (1-q_P)^{T-1-k}$ are the probabilities of observing k public signal realizations equal to 1, conditional on $V = 0$ and $V = 1$, respectively. Under the parameters of the experiment, $B_0^W(p_1) = B_1^W(p_1) \approx -0.075 < 0$. Therefore, applying Lemma A3 part 1), in any equilibrium all traders trade immediately, proving Proposition 2 part a).

Proposition 2, part b). As with $B_x^W(p_1) < 0$, $B_x^R(p_1) > 0$ can be satisfied for many combinations of parameters. To evaluate $B_x^R(p_1)$ for the parameters of treatment W ($\underline{q} = \frac{13}{24}$, $\overline{q} = 1$, $q^* = \frac{17}{24}$, $p_1 = \frac{1}{2}$, and $T = 8$), I first rewrite the benefit when $n - 1$ trades and $T - 1$ public signals occur while waiting as

$$B_x^R(p_1) = \frac{p_1(1-p_1)}{Pr(\underline{s}_i = x)} \left[\sum_{k=0}^{T-1} \sum_{j=0}^{n-1} \frac{C_0(k)C_1(k)D_0(j)D_1(j)(2\overline{q}-1)}{p_1 C_1(k)D_1(j) + (1-p_1)C_0(k)D_0(j)} - (2\underline{q}-1) \right] \tag{6}$$

where $D_0(k) = \frac{(n-1)!}{j!(n-1-j)!}(1-\underline{q})^j \underline{q}^{n-1-j}$ and $D_1(j) = \frac{(n-1)!}{j!(n-1-j)!}\underline{q}^j(1-\underline{q})^{n-1-j}$ are the probabilities of observing j buys by the $n-1$ other traders, conditional on $V=0$ and $V=1$ respectively, and $C_0(k)$, $C_1(k)$ are as above. Under the parameters of the experiment, $B_0^R(p_1) = B_1^R(p_1) \approx 0.128 > 0$. Therefore, applying Lemma A3 part 2), in the unique equilibrium, all traders wait to trade until $t=T$, establishing Proposition 2 part b).

A3. Diagnostic Model

In this section, I establish Proposition D2 for the Diagnostic model (Proposition D1 was proven in Section A1). The payoff-relevant state in this model is the price in the market at the time a trader arrives and whether one's predecessor rushed or waited. If she waited, she'll trade before one gets a second opportunity to trade, making waiting costly.

In treatment DR, where $\overline{q} < \underline{q}$, inspection of the benefit function with $q = \underline{q}$ when no informative events occur shows that it is identically zero. Intuitively, because the second signal never changes one's trading decision, it is of no value so, as long as no cost of information exists, one is indifferent between rushing and waiting. If, however, another trader trades, then Proposition 1 of Kendall (2015) applies, so that the benefit function is strictly negative. Therefore, the unique equilibrium of treatment DR is for every trader except the last to rush independent of the timing decision of the previous trader and the price in the market at the time of arrival. If a trader instead waits, her successor rushes to avoid a strictly negative benefit. But this means that the trader herself faces a strictly negative benefit from waiting. The trader that arrives at T is an exception because she can wait with impunity if her predecessor rushes. In this case, she is indifferent between waiting and rushing.

For treatment DW, I establish that t's benefit from waiting is positive at all p_t, no matter what the strategies of the other traders are. From Proposition 1 of Kendall (2015) if we know that the benefit is positive when $t-1$ waits and $t+1$ rushes, then we know that it must be positive for all combinations of strategies because they result in less potential trades while t waits. Lemma C1 of Kendall (2015) establishes that a benefit function that is positive at $p_t = \frac{1}{2}$ is positive at all prices. Numerically evaluating the benefit function when $t-1$ waits and $t+1$ rushes by plugging the appropriate events into (4) and using $p_t = \frac{1}{2}$, $\underline{q} = \frac{13}{24}$, and $\overline{q} = \frac{17}{24}$ gives $0.130 > 0$. Thus, traders in DW must wait at every history, establishing the second part of Proposition B2.

B. Instructions

I provide the instructions for the W treatment, followed by those for the DW treatment. The instructions for the R and DR treatments are identical except for the differences in parameters.

Instructions

This is a research experiment designed to understand how people make stock trading decisions in a simple trading environment. To assist with our research, we would greatly appreciate your full attention during the experiment. Please do not communicate with other participants in any way and please raise your hand if you have a question.

You will participate in a series of 30 **trials**. In each trial, you and the 7 other participants will each trade a stock with the computer. You will be paid $5.00 for completing all trials. In addition, in each trial you will earn lottery tickets as described next. The more lottery tickets you have, the more you will earn on average.

Trading and Profits

Before each trial, the computer will randomly select whether the stock's value, V, is 100 tickets or 0 tickets. Each is equally likely to be selected. At the start of each trial, you will be given 100 tickets. You may choose either to buy or sell the stock in any one of 8 trading periods. You can only trade **ONCE** and **MUST** trade in one of the periods. If you buy the stock, you will gain its actual value minus the price, P, you pay for it. If you want to sell the stock, you must first 'borrow' it from the computer and later pay back its actual value. So, if you sell, you will gain the price you receive minus its value. In summary, your total profit is:

100+V-P lottery tickets if you buy

100+P-V lottery tickets if you sell

So, for example, if you sell the stock at a price of 50 and it turns out to be worth 100, you would earn 100+(50-100)=50 tickets for that trial. But if it turns out to be worth 0, you would earn 100+(50-0)=150 tickets.

Clues about Value

To help you guess the value of the stock the computer chose, you will get a single **Private Clue** at the start of the trial. This Private Clue will be known **only by you**. Specifically, there will be two possible bins: one that the computer will use if the value is 100 and another that the computer will use if the value is 0. Bins contain some number of blue and green marbles as shown below. The computer will draw a marble randomly from the bin and show it to you as your Private Clue. Each marble in the bin is equally likely to be drawn. The color of the ball you see can give you a hint as to the stock's value.

After observing your Private Clue, you can choose to trade immediately in the first trading period. Alternatively, you can choose to wait and trade in one of the following 7 trading periods. Between trading periods, there are public announcement periods. In each annoucement period, a **Public Clue** will become available. The Public Clue, unlike your Private Clue, is seen by **everyone**. For the Public Clue,the computer will draw a marble from another bin. The bin used will again depend on the stock's true value but the possible bins used for the Public Clue are different from the bins used for the Private Clue, as shown below. Note that the bins used for both clues are fixed throughout the trial - they depend only on the initially chosen random value of the stock. Also, marbles for both types of clues are always replaced before another is drawn.

	Contents of bin if value = 100	Contents of bin if value = 0
Private Clue	13 blue 11 green	11 blue 13 green
Public Clue	17 blue 7 green	7 blue 17 green

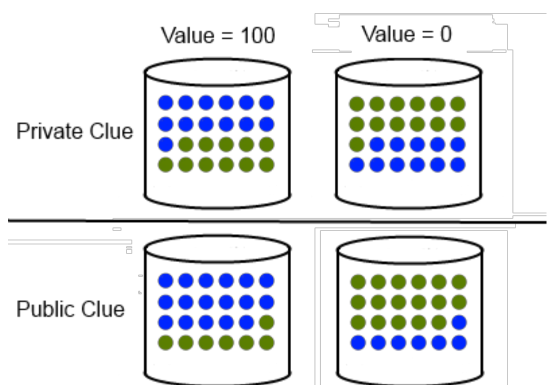

If you decide to wait to trade until the last trading period, the true value of the asset will be revealed to you **before** you trade. Otherwise, you will only learn the true asset value after the trial is complete. Importantly, however, each time you choose to wait, the price is likely to change before your next chance to trade, as described next.

Prices

The price of the stock is set by a Price Setter played by the computer. The Price Setter's job is to set the price equal to the stock's mathematical **expected value** given all of the public information available. Therefore, the price will always be between 0 and 100. The Price Setter can observe the trades made by you and the other participants and the Public Clues. However, she **can not** observe any of the Private Clues nor the true asset value (even when it is revealed in the final period).

Because the price changes with the available information, if you decide to wait, the price at which you can trade is likely to change. The initial price of the stock is 50 tickets, reflecting the fact that it is equally likely to be worth 100 tickets or 0 tickets. After a participant buys, the Price Setter will increase the price and after a participant sells, she'll decrease the price. After a Public Clue is revealed, the price will increase if it suggests (on its own) that the stock is more likely to be worth 100 and decrease if it suggests it is more likely to be worth 0.

Trading Screen

The trading screen you will use to trade is as shown below. The eight trading periods are indicated by the numbers 1-8. Public Clues are revealed between trading periods at the times indicated by the megaphone symbol (the clues are not shown in the figure - they appear elsewhere on your trading screen). All past trading decisions and prices are displayed. Trading periods in which one or more trades occur are indicated by a solid dot. The number of buys is indicated by a "+" and then a number and the number of sells by a "-" and then a number. The current price at which you can trade (74.16 in ths example) is displayed at the current time which is indicated by the dashed red line. The dashed red line will progress to the right as we move through the periods.

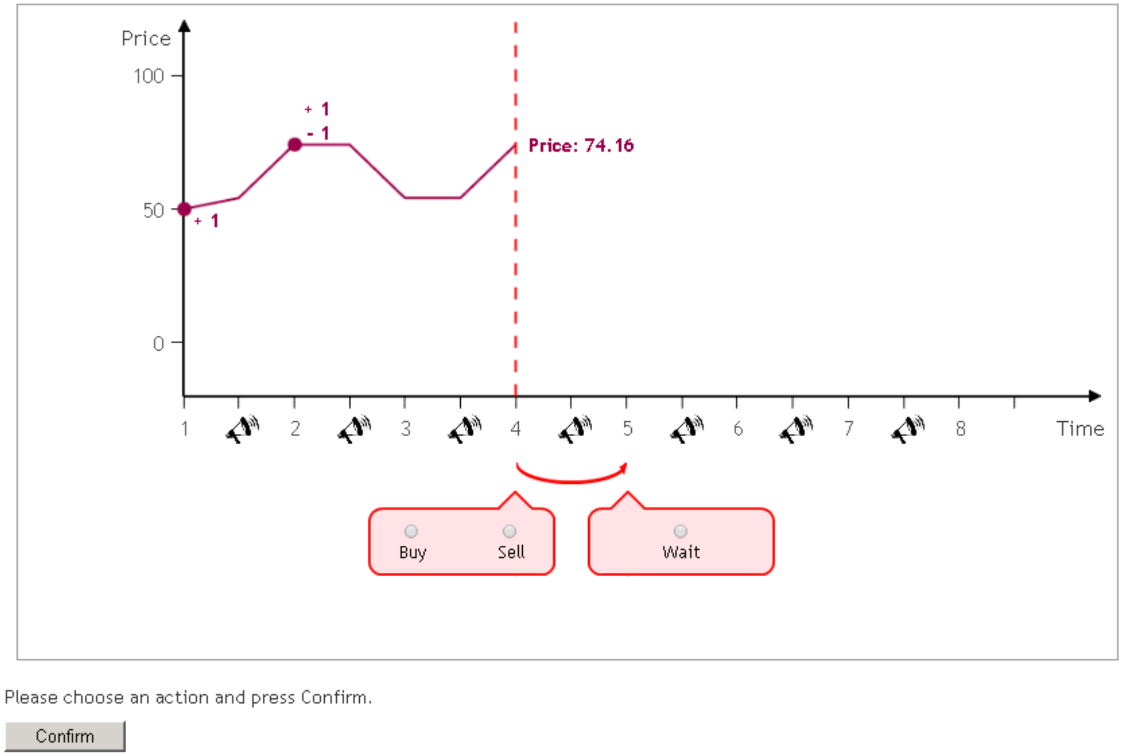

In this example, one participant bought in the first trading period, so the price increased. In the second trading period, one participant bought and one sold, so the price did not change. In the third trading period, no one traded. The first and third Public Clues suggested the stock's value is 100 and the second that the stock's value is 0. It is currently the fourth trading period. Note that this is an example only and is not meant to suggest when you should trade.

In each trading period in which you haven't already traded, you must choose buy or sell or wait and then press the 'confirm' button. If you choose to wait, the red arc points to the next trading period in which you can trade. In trading periods after you have traded, you do not have to do anything - you will simply be notified that you have already traded. In the periods with Public Clues, you must press "OK" to acknowledge having seen the clue.

Summary

1. At the beginning of each trial, the computer randomly selects the stock's value: 0 or 100.
2. Each participant is shown their Private Clue from **the same bin**. Marbles are replaced so each participant may see the same (or different) marbles.
3. Each participant chooses to buy, sell, or wait in the first trading period. After all participants have made their decisions, a Public Clue is revealed and we move to the second trading period.
4. Step 3 is repeated until all 8 trading periods are complete. You must trade in one of the eight trading periods and may only trade once.
5. Just before the 8th trading period, the true value of the stock will be revealed to you if you have not already traded.
6. At each point in time, all past prices and trading decisions are available to use to help guess the value of the stock (in addition to the Clues).

After all participants have traded, the trial is complete. The true value of the asset will be revealed to all participants and you will be told how many tickets you earned for the trial. You will press 'next trial' to participate in the next trial. It is important to remember that, in each trial, the value of the stock is independently randomly selected by the computer -- there is no relationship between the value selected in one trial and another.

After all trials are complete, one lottery will be conducted for each trial. For each lottery, a random number less than 200 will be chosen by the computer. If the number is smaller than the number of lottery tickets you earned for the trial, you will get $1.00. Therefore, the more lottery tickets you earn in each trial, the more you can expect to make (partial tickets are possible and count as well). For example, if you earn 100 tickets in each trial, you can expect to make 0.5*30*1=$15.00 over the 30 trials. But, if you earn 130 tickets in each trial, you can expect to make 0.65*30*1=$19.50.

Please try to make each trading period decision within 15 seconds so that the experiment can finish on time. A timer counts down from 15 to help you keep track of time. Note, however, that if the timer hits zero, you can still enter your trading or wait decision and will still have the same chance to earn money. Before beginning the paid trials, we will have two practice trials for which you will not be paid. These trials are otherwise identical to the paid trials.

Quiz

Please answer the following questions and press the 'Check answers' button to see whether or not you answered all questions correctly. To ensure all participants understand the instructions, everyone must answer all of the questions correctly before we begin the experiment.

---

1. Your Private Clue is a blue marble. Based only on this information, What is the most likely value of the stock?

○ 100

- 0

2. It is the second trading period. You observe another participant sold the stock in the first trading period. What color marble is their Private Clue most likely to be?

- green
- blue

3. Your Private Clue is a blue marble. What color marble is another participant's Private Clue likely to be?

- green
- blue

4. If you choose to sell the stock at a price of 80 and its value turns out to be 100, how many total tickets would you get for that trial?

- 80
- 20
- 180

5. If you choose to buy the stock at a price of 25 and its value turns out to be 100, how many total tickets would you get for that trial?

- 25
- 75
- 175

6. The stock's current price is 80. Which value of the stock is more likely?

- 100
- 0

Once you have completed the quiz, please press 'Check answers'.

Check answers

Instructions

This is a research experiment designed to understand how people make stock trading decisions in a simple trading environment. To assist with our research, we would greatly appreciate your full attention during the experiment. Please do not communicate with other participants in any way and please raise your hand if you have a question.

You will participate in a series of 42 **trials**. In each trial, you and 5 other people will each trade a stock with the computer. The trades are sequenced so that each person arrives to the market one at a time. You will be paid $5.00 for completing all trials. In addition, in each trial you will earn lottery tickets as described next. The more lottery tickets you have, the more you may earn.

Before each trial, the computer will randomly select whether the stock's value, V, is 100 tickets or 0 tickets, with equal probability. When it is your turn, you will be given 100 tickets to begin and then you may choose either to buy or sell the stock. You will only trade **once**. If you buy the stock, you will gain its actual value minus the price, P, you pay for it. If you want to sell the stock, you must first 'borrow' it from the computer and later pay back its actual value. So, if you sell, you will gain the price you receive minus its value. In summary, your total profit is:

100+V-P lottery tickets if you buy

100+P-V lottery tickets if you sell

So, for example, if you sell the stock at a price of 50 and it turns out to be worth 100, you would earn 100+(50-100)=50 tickets for that trial. But if it turns out to be worth 0, you would earn 100+(50-0)=150 tickets.

To help you guess the value of the stock the computer chose, you will get a First Clue when it is your turn. Specifically, there will be two possible bins: one that the computer will use if the value is 100 and another that the computer will use if the value is 0. Bins contain some number of blue and green marbles as shown below. The computer will draw a marble randomly from the bin and show it to you. Each marble in the bin is equally likely to be drawn. You can use this clue to get a better idea of what the stock's value is.

After observing your First Clue, you can choose to buy or sell the stock immediately. Alternatively, you can choose to wait and trade in the following period. If you choose to wait, you will get a Second Clue: the computer will draw a marble from another bin. The bin used will again depend on the stock's true value but the possible bins used for the Second Clue are different from the bins used for the First Clue as shown below. After observing the Second Clue, you may choose to buy or sell the stock. Note that you must trade in either the period you arrive or the next period and that you can **only trade once**.

	Contents of bin if value = 100	Contents of bin if value = 0
First Clue	13 blue 11 green	11 blue 13 green
Second Clue	17 blue 7 green	7 blue 17 green

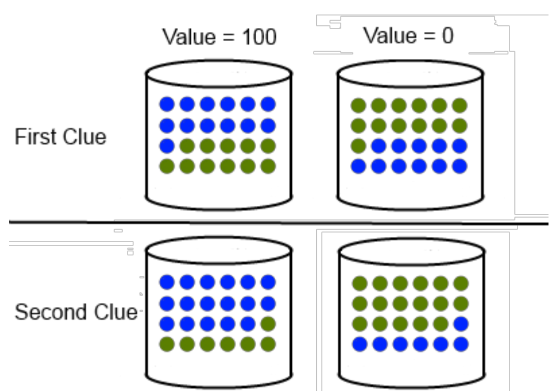

Importantly, if you decide to wait, other participants will have a chance to trade before you have your second chance to trade. Specifically, if the participant before you waited, they will trade **before** you. And, the next participant will arrive to the market and have a chance to trade immediately or to wait, just like you. If they trade immediately, they will also trade **before** you. The timing of trades is shown in the figures below. Each open circle indicates a possible trade time and the arcs indicate the two possible times at which a particular participant can trade. The red, highlighted arc corresponds to the participant whose turn it is (the second participant in this example).

When you arrive in the market, you will observe whether or not the person before you has already traded. In the first example to the right, the second participant has just arrived to the market and the first participant has already traded (they didn't wait). This is indicated by the arrow pointing to the time at which they traded. Also, the open circle contains a '+' sign (for a buy) or a '-' sign (for a sell) after a trade has completed. In this example, if the second participant waits, there is only one possible trade (by the third participant if they trade immediately) before the second participant gets a chance to trade again.

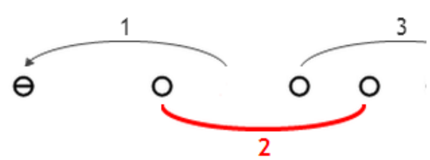

In the second example to the right, when the second participant arrived in the market, she saw that the first participant has waited (arrow to the right). In this case, the open circle where the first participant arrived disappears because no trade occurred. Also, a shaded circle appears where the first participant must trade. So, if the second participant were to wait, the first participant would trade for certain before the second participant has a chance to trade again. And, the third participant would also trade before the second, if they trade immediately.

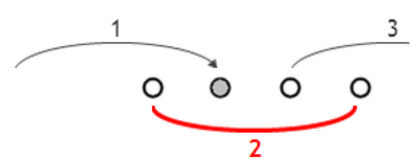

All past trading decisions and prices will be displayed in a figure as shown below. In this example, the first participant waited and then bought the stock. The second participant sold the stock immediately. The third participant has waited and it is currently the fourth participant's turn. The price history is displayed in the graph along with the current price at which you can trade (67.27 in ths example). The current time period is indicated by the dashed red line. After choosing buy or sell (or wait in the period you arrive), you will press the 'confirm' button (not shown).

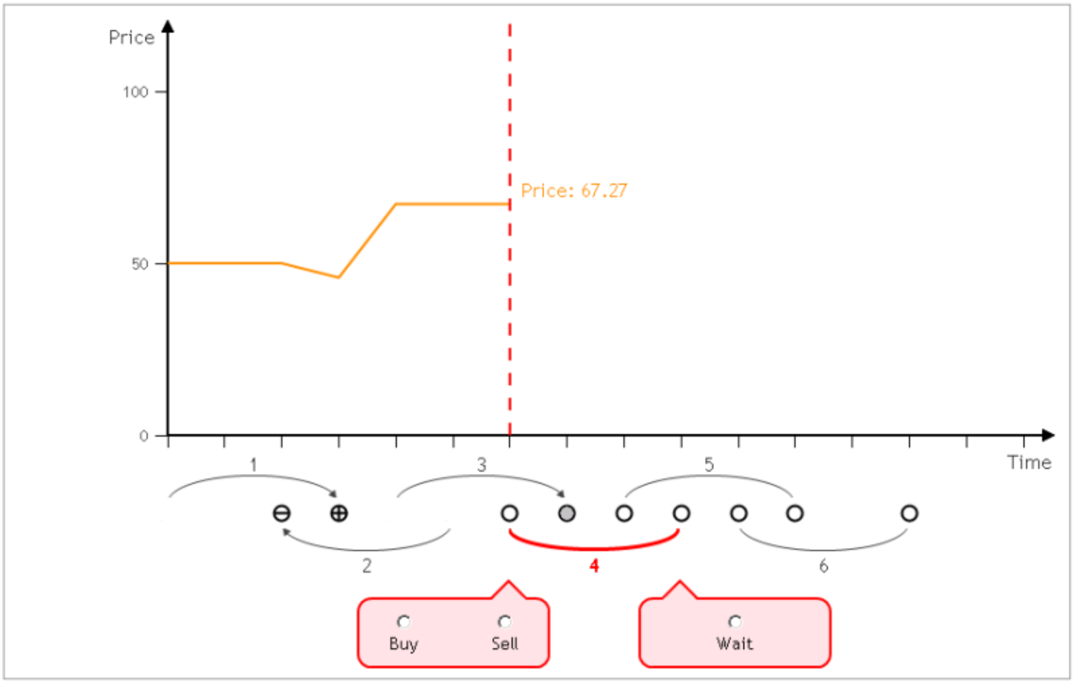

The price of the stock will change over time based upon the trades that occur. Therefore, if you wait, and another participant trades, the price you will face when you have a second chance to trade will be **different** than the price you could have traded at if you traded immediately.

The initial price of the stock is 50 tickets, reflecting the fact that it is equally likely to be worth 100 tickets or 0 tickets. Similarly, after a trade, the price of the stock will be updated so that it continues to be equal to the stock's expected value given the information that can be inferred from the trades, but no other information. The stock's price will therefore increase after a buy and decrease after a sell, and the change will be smaller for immediate trades than for trades that take place after waiting (as you can see in the example above).

Summary

1. At the beginning of each trial, the computer randomly selects the stock's value: 0 or 100.
2. The first participant is shown their First Clue and then chooses to buy, sell, or wait.
3. If the first participant chooses to buy or sell, they are done for this trial. If they choose to wait, the second participant arrives to the market, observes their own First Clue, and chooses to buy, sell, or wait.
4. If the first participant chose to wait, they will then buy or sell after observing a Second Clue.
5. Steps 2-4 are repeated for all six participants. Each participant observes their own First Clue (and Second Clue if they choose to wait) from **the same bins** as previous participants. Marbles are always replaced so later participants may see the same (or different) marbles.
6. At each point in time, all past prices and trading decisions are available to use to help guess the value of the stock (in addition to the Clues).

After all participants have traded, the trial is complete. The true value of the stock will then be revealed and you will be told how many tickets you earned for the trial. You will press 'next trial' to participate in the next trial. It is important to remember that, in each trial, the value of the stock is independently randomly selected by the computer -- there is no relationship between the value selected in one trial and another. The order of participants in the sequence may be different from trial to trial: you may be participant 1 in one trial, participant 4 in another, and participant 6 in yet another. Which participant you are in each trial will be told to you before you make your trade.

After all 42 trials are complete, one lottery will be conducted for each trial. For each lottery, a random number less than 200 will be chosen by the computer. If the number is smaller than the number of lottery tickets you earned for the trial, you will get $1.00. Therefore, the more lottery tickets you earn in each trial, the more you can expect to make (partial tickets are possible and count too). For example, if you earn 100 tickets in each trial, you can expect to make 0.5*42*1=$21.00 over the 42 trials. But, if you earn 120 tickets in each trial, you can expect to make 0.6*42*1=$25.20.

Please try to make each decision within 15 seconds so that the experiment can finish on time. A timer counts down from 15 to help you keep track of time when it is your turn. Note, however, that if the timer hits zero, you can still enter your trading decision and will still have the same chance to earn money. Before beginning the paid trials, we will have two practice trials for which you will not be paid. These trials are otherwise identical to the paid trials except that they are not timed.

Quiz

Please answer the following questions. To ensure you understand the instructions, you must answer all of the questions correctly before we begin the experiment.

1. You are the first participant and have waited. Your First Clue is a green marble and your Second Clue is a blue marble. Based only on this information, What is the most likely value of the stock?

- 100
- 0

2. You are the second participant and observe the first participant

- green

sold the stock immediately. What color marble are they most likely to have seen?

- blue

3. You are the first participant and your First Clue is a blue marble. What color marble is the second participant most likely to see?

- green
- blue

4. If you choose to sell the stock at a price of 80 and its value turns out to be 100, how many total tickets would you get for that trial?

- 80
- 20
- 180

5. If you choose to buy the stock at a price of 25 and its value turns out to be 100, how many total tickets would you get for that trial?

- 25
- 75
- 175

6. The stock's current price is 80. Which value of the stock is more likely?

- 100
- 0

7. You are the second participant and you observe that the first participant did not trade immediately. How many trades can occur before you trade again if you wait?

- 0
- 1 or 2
- 0 or 1
- 1

Once you have completed the quiz, please press 'Check answers'.

Check answers