

Herding and Contrarianism: A Matter of Preference?

Chad Kendall*

February 5, 2021

Abstract

Herding and contrarian strategies produce informational inefficiencies when investors ignore private information, instead following or bucking past trends. In a simple market model, I show theoretically that investors with prospect theory preferences generically follow herding or contrarian strategies, but do so because of future returns as opposed to past trends. I conduct a laboratory experiment to test the theory and to obtain an estimate of the distribution of preferences in the subject population. I find that approximately 70% of subjects have preferences that induce herding. Using the preference estimates, I quantify informational efficiencies and predict trade behavior in more general environments.

1 Introduction

Herding and momentum strategies, and their antithesis, contrarian strategies, interest financial economists because of their implications for the informational efficiency of market prices, which are in turn important due to potential effects on the real economy (see Bond, Edmans, and Goldstein (2012) for a recent review).¹ What drives investors to use these strategies and

*Department of Finance and Business Economics, Marshall School of Business, University of Southern California, 701 Exposition Blvd, Ste. 231 HOH-231, MC-1422, Los Angeles, CA 90089-1422 (e-mail chadkend@marshall.usc.edu). Supplementary material is available in the Online Appendix here: <https://www.sites.google.com/site/chadwkendall/home/files/kendall-ptherding-appendix.pdf>. I would like to thank Kenneth Ahern, Nick Barberis, Wayne Ferson, Cary Frydman, Valentin Haddad, Larry Harris, Lawrence Jin, Ryan Oprea, Fernando Zapatero, and various seminar participants for valuable feedback. JEL codes: G40, G14, D90

¹By a ‘momentum strategy’, I’m referring to the strategy of buying if prices have risen and selling if they’ve fallen, sometimes referred to as positive feedback trading (De Long (1990)), and not to be confused with the cross-sectional strategy of shorting losers and going long winners (which goes by the same name). Momentum and contrarian strategies are frequently discussed by practitioners, and a small academic literature suggests different types of individuals or firms pursue each (see Grinblatt, Titman, and Wermers (1995), Grinblatt and Keloharju (2000), Baltzer, Jank, and Smajlbegovic (2019), and Grinblatt et al. (2016)).

to choose one over the other? The most basic explanation for herding/momentum (following the trend) is simply imitation, a form of ‘herd mentality’ (Mackay (1841)), but behavioral explanations such as extrapolative expectations (De Long et al. (1990), Greenwood and Shleifer (2014), Barberis, Greenwood, Jin, and Shleifer (2018)) and rational explanations based on information externalities (Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)) have also been theorized.² Similarly, contrarian strategies (going against the trend) have rational explanations given a suitable informational environment (Avery and Zemsky (1998) and Park and Sabourian (2011)). In all of these explanations, the observation of previous investor actions and/or prices is critical and, in fact, the strategies are often defined in terms of inference from past trends (Avery and Zemsky (1998) and Park and Sabourian (2011)). In this paper, I demonstrate an additional channel through which past trends can appear to produce herding and contrarian strategies, one in which inference and/or imitation is not necessary. Rather, due to the nature of the relationship between past price trends and subsequent expected returns, these strategies may simply be driven by preferences over future returns.

I make my case for a preference-based explanation using both theory and a laboratory experiment. In the theoretical contribution, I revisit a standard model of trading with asymmetric information, that of Glosten and Milgrom (1985). In this model, herding and contrarianism are impossible with risk-neutral investors (Avery and Zemsky (1998)), but both behaviors occur in past experimental tests of the model (Cipriani and Guarino (2005,2009) and Drehmann, Oechssler, and Roider (2005)).³ To resolve this discrepancy, I consider the behavior of an investor with cumulative prospect theory (CPT) preferences (Kahneman and Tversky (1992)). I show that such an investor generically (for almost all preference parameters) either buys or sells the tradable asset at extreme prices, *independently* of her private information, thus reducing the informational efficiency of the market by not revealing her signal. These behaviors do not, however, rely on an investor making inferences from the behavior of others or changes in prices. Instead, the key fact is that extreme prices imply highly skewed forward-looking returns: at high prices, returns tend to be negatively skewed, and conversely, at low prices, returns tend to be positively skewed. If investors then have strong preferences for skewness, it dominates their decisions, overwhelming their private information.

The key components of CPT, probability weighting and an S-shaped value (utility) func-

²Momentum and herding produce similar trade patterns, but herding is often defined in the context of asymmetric information environments. See Devenow and Welch (1996) and Hirshleifer and Teoh (2003) for surveys of herding in financial markets.

³Both experimental papers show that believing that previous subjects make mistakes can justify contrarian behavior, but neither paper provides an explanation for the herding behavior observed.

tion, provide such strong skewness preferences. As has been well-recognized (Barberis and Huang (2008)), the overweighting of small probabilities in CPT generates a preference for positive skewness: investors overweight the small probability that the asset will generate a large return. Less obviously, an S-shaped value function generates a preference for negative skewness due to risk-aversion over positive returns and risk-seeking over negative returns. Unless these opposing effects exactly offset each other, an investor has preferences for either positive or negative skewness. A preference for negative skewness (value function curvature dominates probability weighting), leads an investor to herd, buying at high prices when returns are negatively skewed and taking the opposite position (selling) at low prices when returns are positively skewed. A preference for positive skewness generates the opposite: contrarian behavior in which the investor sells at high prices, and buys at low prices. In either case, investor behavior will appear to depend on past trends because high prices tend to follow price increases, and conversely.

In the theoretical model, investors arrive sequentially to a market, trading a single, binary-valued asset with a market maker who posts separate bid and ask prices. Each investor, after receiving a private, binary signal about the asset's fundamental value, may buy or sell a single unit of the asset (or abstain from trading). With a binary-valued asset, high prices immediately imply that the expected return from buying the asset is negatively skewed: it is very likely to return a small positive amount (when the asset is valuable), but occasionally results in a large loss (when it is not). Selling the asset at high prices instead results in a positively skewed return. The standard result with risk-neutral, expected utility investors (Avery and Zemsky (1998)) is that each investor trades according to her private information - buying with a favorable signal and selling otherwise. With CPT preferences, I show that an investor ignores her private information at extreme prices, instead trading in a single direction. Consistent with the previous intuition, the direction of her preferred trade depends upon the relative strengths of value function curvature and probability weighting. In fact, with power utility and a single parameter probability weighting function, the direction of trade depends simply on the difference of the two parameters that govern these functions. Thus, CPT preferences provide a unifying explanation for herding and contrarianism.

My theoretical model predicts that CPT preferences *can* produce herding and contrarianism, but to what extent *do* they? Disentangling a preference-based explanation from social-learning explanations would be very difficult in naturally-occurring data where preferences and information are typically unobservable. Even the data from past laboratory experiments (Cipriani and Guarino (2005,2009) and Drehmann, Oechssler, and Roider (2005)) is difficult to interpret because preferences, belief-based errors, as well as beliefs about others' preferences and errors all contribute to behavior. I therefore conduct a new laboratory experiment

in which I isolate the role of preferences. Critically, I shut down all social causes of behavior by having subjects make decisions in an environment in which their decisions are independent of others'. However, I maintain the framing of a market: subjects observe a past price trend and make buy or sell decisions. Doing so allows me to observe decisions when choices are framed in a natural way and to test for any direct effect of trends on decision-making.

In the first (MAIN) treatment, I control for belief-based errors (such as extrapolative expectations) by providing subjects with the correct Bayesian posteriors corresponding to their information sets (including both public and private information), allowing me to isolate the effect of preferences. By having subjects make a large number of trading decisions (60), I am able to estimate the distribution of preferences in the population and use it to assess (via simulation) the informational efficiency of markets composed of this population. In the second (NO BELIEFS) treatment, I determine whether belief formation errors exacerbate herding and contrarian behavior by keeping everything identical to the MAIN treatment except for the provision of correct Bayesian posteriors.

I find strong evidence of herding and contrarianism in both treatments, clearly demonstrating that social learning is not necessary to generate these behaviors. Moreover, comparing the two treatments, I find that, contrary to the prediction of extrapolative expectations, belief-based errors actually *reduce* herding while leaving the level of contrarian behavior unchanged. In the MAIN treatment, where these errors are absent, approximately three-quarters of subjects exhibit either herding or contrarian behavior and, among these, 90% follow herding strategies. I find only weak evidence that the precise nature of the trend (i.e. whether it is monotonic or not) affects decisions, suggesting that, consistent with the theory, it is the skewness of future returns rather than the price trends themselves that drive behavior.

Estimating preference parameters using data from the MAIN treatment, I find that 75% of subjects are better described by CPT preferences than expected utility (CRRA). Importantly, this result does not depend on CPT having more degrees of freedom (a general critique of prospect theory), given that I took an *ex ante* stance on the reference point (the status quo) and establish theoretically that (under the assumed functional forms) a simple function of the difference between the degree of value function curvature and probability weighting determines behavior. In fact, CPT better describes behavior than expected utility even when CPT is restricted to a single degree of freedom. Of the minority of subjects that are not best described by CPT, most are risk-neutral expected utility maximizers, but a few abstain from trading at extreme prices. This latter behavior cannot be reconciled with CPT preferences (but is consistent with a high level of risk-aversion), illustrating that, while the theory is consistent with the behavior of most subjects, it is falsifiable.

I perform two exercises with the estimated distribution of preferences. First, I simulate markets composed of a sequence of investors (as in the theory) to compare informational efficiency relative to a benchmark market comprised of only risk-neutral investors. The simulations demonstrate not only that herding and contrarianism produce significant informational inefficiencies relative to the benchmark, but also that information cascades can develop, leading to stagnate prices. Second, I evaluate the types of Bayesian errors subjects make, on average, assuming that the distribution of preferences is the same across treatments. At intermediate priors, they underweight private information, while at extreme priors they underweight the information contained in the public signals, consistent with a large literature on underinference (Benjamin (2018)).

Lastly, I take the estimated preferences of the modal subject in the experiment (who exhibits herding behavior) and ask how an investor with these preferences would behave when facing actual market returns. Because market returns become more (less) skewed as prices rise (fall), a fact first documented by Chen, Hong and Stein (2011), if preferences over skewness drive behavior as in the theory and experiment, we should expect this investor to be willing to purchase an asset after rising prices even with negative private information, and to sell with positive information after falling prices. Simulation results confirm this conjecture, demonstrating that preferences can drive herding in environments more general than the simple, binary asset value setup of the theory and experiment.

The seminal paper on herding in financial markets is that of Avery and Zemsky (1998) who show that herding and contrarianism are impossible with risk-neutral investors (unless additional sources of uncertainty are added to the model). More recent theoretical papers have studied whether or not non-expected utility preferences can generate these behaviors. Qin (2015) shows that regret aversion generates herding, but not contrarian, behavior. Boortz (2016) builds on Ford, Kelsey, and Pang (2013), to show that ambiguity can generate both behaviors, but only with preferences that vary with the state. In contrast to these models, I show that static CPT preferences generate both types of behavior. Other papers have considered the possibility that hedging motives (Decamps and Lovo, 2008) or private values (Cipriani and Guarino, 2008) can induce herd behavior. The model here, instead, explains how herding and contrarianism can arise under common values and without hedging motives - in an environment identical to that of previous experiments in which these behaviors have been observed.

The paper also contributes to the experimental literature on herding in financial markets (see Cipriani and Guarino (2005), Drehmann, Oechssler, and Roider (2005), Cipriani and Guarino (2008), Cipriani and Guarino (2009), Park and SgROI (2012), and Bisiere, Decamps, and Lovo (2015)). Bisiere, Decamps, and Lovo (2015) introduced the key methodological

innovation of converting the market into an individual decision problem in order to study how preferences affect the revelation of private information. However, their study differs in that they focus on an environment in which hedging motives can cause standard risk preferences to generate herding and contrarianism. I illustrate these behaviors even in an environment without hedging motives and put forward prospect theory as an explanation, a theory they do not consider.⁴ More broadly, this paper contributes to a growing literature that applies CPT preferences to understanding investor behavior in financial markets. See Barberis (2013) for a review.

2 Theory

2.1 Model

The model is a sequential trading model based on that of Glosten and Milgrom (1985). In each period $t = 1, 2, \dots, T$, a single new investor arrives to the market to trade an asset of unknown value, $V \in \{0, 1\}$. I denote the initial prior that the asset is worth 1 by $p_1 \in (0, 1)$. Upon arrival, an investor may either buy or sell short a single unit, or not trade, $a_t \in \{buy, sell, NT\}$. After making her decision, the investor leaves the market. All trades are with a risk-neutral market maker who is assumed to face perfect competition, earning zero profit in expectation. The market maker incorporates the equilibrium informational content of the current order in setting prices. Specifically, he posts an ask price, A_t , at which he is willing to sell a unit of stock and a bid price, B_t , at which he is willing to buy a unit. When the asset value is realized at T , investors who purchased the asset at time t receive a payoff of $V - A_t$ and those who sold receive a payoff of $B_t - V$ (no discounting). All market participants observe the complete history of trades and prices, denoted $H_t = \{a_1, a_2, \dots, a_{t-1}\} \cup \{A_1, A_2, \dots, A_{t-1}\} \cup \{B_1, B_2, \dots, B_{t-1}\}$.

Investors are one of three types: risk-neutral, prospect theory, or uninformed investors. Uninformed investors, who arrive with probability $1 - \mu$, $\mu \in (0, 1)$, trade for exogenous reasons and are equally likely to buy or sell (such traders are sometimes referred to as ‘noise’ traders). Risk-neutral investors have standard risk-neutral expected utility preferences and arrive with probability, $\mu\gamma$, $\gamma \in (0, 1)$. Finally, prospect theory investors have the CPT preferences of Kahneman and Tversky (1992) and arrive with the remaining probability, $\mu(1 - \gamma)$.

CPT specifies value functions, $v^+(\cdot)$ and $v^-(\cdot)$, and decision weight functions, $w^+(\cdot)$ and

⁴In two robustness treatments, Bisiere, Decamps, and Lovo (2015)) eliminate the hedging motive. They continue to observe herding and contrarianism, inconsistent with risk-averse expected utility maximization but consistent with the model here.

$w^-()$, over gains and losses, respectively, with respect to a reference point, which I take to be the status quo. The decision weight functions apply to capacities, a generalization of probabilities, but for binary outcomes result in simple non-linear transformations of the objective probabilities. Specifically, the utility a CPT investor derives from a binary lottery, $\mathcal{L}$, which returns a gain of x with probability r and a loss of y with probability $1 - r$ is given by $U(\mathcal{L}) = w^+(r)v^+(x) + w^-(1 - r)v^-(y)$.

I adopt the functional forms for the value and decision weight functions from the original work of Kahneman and Tversky (1992), because they are tractable, parsimonious, and appear to fit decisions over binary gambles reasonably well. Specifically, I assume $v^+(x) = x^\alpha$, $v^-(y) = -\lambda(-y)^\alpha$, and $w^+(r) = w^-(r) = \frac{r^\delta}{(r^\delta + (1-r)^\delta)^{\frac{1}{\delta}}}$, with $\alpha \in (0, 1]$, $\lambda \geq 1$, and $\delta \in (0, 1]$. $\alpha \in (0, 1)$ reflects the common experimental finding of risk-aversion over gains and risk-seeking over losses (an “S-shaped” value function). $\lambda \geq 1$ reflects loss-aversion: losses are weighted more heavily than gains. Finally, $\delta \in (0, 1)$ implies that low-probability events are overweighted.

Completing the model, I endow the risk-neutral and CPT investors with private information upon arrival to the market: each receives a binary signal, $s_t \in \{0, 1\}$, which has the correct realization with probability $q = Pr(s_t = 1|V = 1) = Pr(s_t = 0|V = 0) \in (\frac{1}{2}, 1)$. All signals are independent conditional on V . I refer to $s_t = 1$ as a *favorable* signal, and $s_t = 0$ as *unfavorable*.

2.2 Solution Concept

Being a game of asymmetric information, the solution concept is Perfect Bayesian Equilibrium. An equilibrium consists of a specification of the strategies of the risk-neutral and CPT investors, along with the bid and ask prices of the market maker. These prices depend upon the market maker’s knowledge of the distribution of preference types as well as his beliefs about their strategies. As usual, these beliefs, which are pinned down at every history due to the presence of the uniformed investors, must be correct in equilibrium. Strategies are functions of the complete history of prices and trades, as well as one’s private signal, to an action: buy, sell, or not trade. As these details are standard, I omit formal definitions.

2.3 Modeling Choices

Two modeling choices were driven by the desire to test the model in the laboratory. First, I adopt the status quo as the reference point. Baillon, Bleichrodt, and Spinu (2019) show experimentally that the status quo and ‘maxmin’ reference points are the most commonly used by subjects among many possible alternatives, including expectations-based reference

points. They define the maxmin reference point as the maximum of the minimum outcomes across gambles in a choice set, which, in the environment here, gives identical predictions to a status quo reference point. The behavioral asset pricing literature has tended to use the expected wealth from investing in a risk-free asset (e.g. Barberis and Huang (2008)) rather than the status quo. But, in the absence of a risk-free asset, as is the case here, these different specifications are equivalent in the sense that the reference point is the amount an investor can attain without risk. Expectations-based reference points, such as those in Kőszegi and Rabin (2006,2007), are another popular alternative. However, I show in Online Appendix H that they are generally inconsistent with the experimental evidence that follows, because they preclude herding.

Second, I include two types of informed investors (risk-neutral and CPT) because we would expect heterogeneous preferences in any population, a hypothesis which I confirm when analyzing the experimental data. Within the CPT investors, I limit the analysis to a single preference type (homogeneous parameters) for ease of exposition only. When I simulate markets based on the estimated distribution of preferences from the experiment, I allow for a distribution of CPT preferences.

2.4 Analysis: Trader Actions

I first establish (Lemma 1 of Online Appendix A) that the equilibrium actions of the risk-neutral investors are the same as in the original work of Glosten and Milgrom (1985). In each period, the market maker posts separate bid and ask prices given by $B_t = Pr(V = 1|H_t, a_t = sell)$ and $A_t = Pr(V = 1|H_t, a_t = buy)$, respectively, a consequence of the assumption of perfect competition. Intuitively, the ask price exceeds the public belief, denoted $p_t = Pr(V = 1|H_t)$, because a buy decision reflects favorable private information, $s_t = 1$, in equilibrium. Similarly, the public belief exceeds the bid price, resulting in the standard bid-ask spread, $A_t - B_t > 0$. Given their private information, risk-neutral investors, buy with favorable signals and sell with unfavorable signals. Intuitively, a favorable signal leads to a private belief greater than the ask price, so that buying leads to an expected profit, and conversely with an unfavorable signal. Importantly, uninformed investors allow the adverse selection problem between informed investors and the uninformed market maker to be overcome. Due to their presence, the bid and ask prices do not fully reflect the private information of informed investors, who are then able to make profitable trades. The market maker loses money to informed investors, but recoups it from uninformed investors.

CPT investor behavior is more novel. Given her utility function, a CPT investor with a private belief, $b_t = Pr(V = 1|H_t, s_t)$, adopts the following optimal decisions:

$$\begin{aligned}
\text{buy if } w^+(b_t)v^+(1-A_t) + w^-(1-b_t)v^-(-A_t) &\geq 0 \\
\text{sell if } w^+(1-b_t)v^+(B_t) + w^-(b_t)v^-(B_t-1) &\geq 0
\end{aligned} \tag{1}$$

where the utility of not trading results in no gain or loss and is normalized to zero. If neither equation in (1) holds, then the CPT investor abstains from trading. Substituting the functional forms for the value and decision-weight functions, results in:

$$\begin{aligned}
\text{buy if } \left(\frac{b_t}{1-b_t}\right)^\delta &\geq \lambda \left(\frac{A_t}{1-A_t}\right)^\alpha \\
\text{sell if } \left(\frac{b_t}{1-b_t}\right)^\delta &\leq \frac{1}{\lambda} \left(\frac{B_t}{1-B_t}\right)^\alpha
\end{aligned} \tag{2}$$

Informed investors have private beliefs (denoted b_t^s) conditional on the history and their private signal, $s_t = 1$ or $s_t = 0$, given by Bayes' rule, $b_t^1 = \frac{p_t q}{p_t q + (1-p_t)(1-q)}$ and $b_t^0 = \frac{p_t(1-q)}{p_t(1-q) + (1-p_t)q}$. Because of perfect competition, Bertrand competition ensures the ask price is the minimum price at which the market maker earns zero expected profit. Similarly, the bid price is the maximum such price, so we have:

$$\begin{aligned}
A_t &= \min \left\{ A_t \in [0, 1] : A_t = \frac{p_t \Pr(a_t=\text{buy}|V=1)}{p_t \Pr(a_t=\text{buy}|V=1) + (1-p_t) \Pr(a_t=\text{buy}|V=0)} \right\} \\
B_t &= \max \left\{ B_t \in [0, 1] : B_t = \frac{p_t \Pr(a_t=\text{sell}|V=1)}{p_t \Pr(a_t=\text{sell}|V=1) + (1-p_t) \Pr(a_t=\text{sell}|V=0)} \right\}
\end{aligned} \tag{3}$$

where the conditional probabilities of observing a purchase or a sale depend upon the proportion of risk-neutral and CPT investors, as well as their equilibrium strategies.

After substituting the expressions for investors' private beliefs and the bid and ask prices, the optimal decision of a CPT investor with a favorable signal becomes

$$\begin{aligned}
\text{buy if } \left(\frac{p_t}{1-p_t}\right)^{\delta-\alpha} &\geq \lambda \left(\frac{1-q}{q}\right)^\delta \left(\frac{\Pr(a_t=\text{buy}|V=1)}{\Pr(a_t=\text{buy}|V=0)}\right)^\alpha \\
\text{sell if } \left(\frac{p_t}{1-p_t}\right)^{\delta-\alpha} &\leq \frac{1}{\lambda} \left(\frac{1-q}{q}\right)^\delta \left(\frac{\Pr(a_t=\text{sell}|V=1)}{\Pr(a_t=\text{sell}|V=0)}\right)^\alpha
\end{aligned} \tag{4}$$

The corresponding equations for an investor with an unfavorable signal are identical except that the ratio of $1-q$ to q on the right-hand side is inverted in each.

Although the opposing effects of α and δ have received relatively little attention in applications of CPT (with the exception of Barberis' (2012) model of casino gambling), it is immediately clear from (4) that their difference, $\delta - \alpha$, plays a critical role. To understand the intuition, consider a simplified example. Assume investors do not receive private signals. In this case, the bid and ask prices collapse to the public belief, p_t , because the market maker no longer needs to infer an investor's private information. Risk-neutral investors in this situation have no incentive to trade because their private beliefs correspond to that of the public belief (equal to price): the gambles corresponding to a purchase or a sale have zero expected value. CPT investors generally have a strict incentive to trade, however. Absent

private information, equations (2) for a CPT investor become

$$\begin{aligned} \text{buy if } & \left(\frac{p_t}{1-p_t} \right)^{\delta-\alpha} \geq \lambda \\ \text{sell if } & \left(\frac{p_t}{1-p_t} \right)^{\delta-\alpha} \leq \frac{1}{\lambda} \end{aligned} \quad (5)$$

We see that, for public beliefs sufficiently different from one-half, either buying or selling is strictly preferred to not trading. $p_t = \frac{1}{2}$ plays a special role in the analysis because of symmetry: an investor with a favorable signal at a public belief, p_t , is in a symmetric situation (faces the same gambles, with buy replacing sell, and vice versa) to an investor with an unfavorable signal at a public belief, $1 - p_t$. Therefore, investor behavior must be symmetric around $p_t = \frac{1}{2}$.

Consider a public belief, $p_t > \frac{1}{2}$. As the decision weights become more distorted from linearity (δ decreases from one), the propensity to buy decreases and the propensity to sell increases. Intuitively, a decrease in δ increases the weight assigned to the small probability, $1 - p_t$, and reduces the weight assigned to the larger probability, p_t . For a purchase, the small probability is associated with a loss and for a sale, it's associated with a gain, so that a decrease in δ makes selling relatively more attractive. Therefore, for p_t sufficiently large and $\delta < \alpha$, the investor strictly prefers to sell the stock: she exhibits a *preference for positive skewness*, the consequence of CPT studied extensively in Barberis and Huang (2008).

Conversely, consider an increase in the curvature of the value function (decrease in α from one). Mathematically, we see that we get exactly the opposite effect from that due to an increase in the distortion of probabilities. Intuitively, as the curvature increases, the small gain ($1 - p_t$) that occurs with probability p_t if one buys is preferred to the large gain (p_t) that occurs with probability $1 - p_t$ if one sells, a simple consequence of risk-aversion (an investor with risk-neutral preferences would be indifferent). At the same time, the small probability of a large loss if one buys is preferred to the large probability of a small loss if one sells, due to risk-seeking. Both effects make buying preferable to selling so that for p_t sufficiently large and $\delta > \alpha$, the investor always strictly prefers to buy: she exhibits a *preference for negative skewness*.

Finally, consider the role of loss aversion. An increase in λ reduces the range of public beliefs at which an investor is willing to trade, because it simultaneously makes each inequality in (5) more difficult to satisfy. The intuition here is simple: an increase in loss-aversion increases the disutility of losses, which makes one more likely to abstain from taking a position in the asset. Perhaps surprisingly, however, loss aversion prevents trading only at intermediate public beliefs. When beliefs become sufficiently extreme, returns become so skewed that trading is preferred to abstention. Intuitively, although the potential losses are larger at extreme public beliefs, one can take the side of the trade that either minimizes the

probability ($\delta < \alpha$) or the size ($\delta > \alpha$) of a loss. At intermediate beliefs, on the other hand, the chance of a loss of medium size and probability can only be avoided by abstaining from trade.

This simple example captures the countervailing forces of distortions due to decision weights and value function curvature. The intuition carries over to the full equilibrium characterization in the following section.

2.5 Analysis: Equilibrium


The equilibrium strategies of investors are functions of the public belief and their private signals. As discussed in the previous section, absent private information, skewness preferences cause CPT investors to strictly prefer to trade in a particular direction as prices become extreme. With private information, because the skewness in returns is unbounded, at sufficiently extreme prices, skewness preferences dominate private information, causing such an investor to trade in the same direction for both realizations of her private signal. Thus, when the public belief is sufficiently large and $\delta > \alpha$, CPT investors buy regardless of their private signal, and, for sufficiently small public beliefs, they always sell. Conversely, when $\delta < \alpha$, CPT investors sell regardless of their private signal when the public belief is sufficiently large, and buy when the public belief is sufficiently small.

For less extreme public beliefs, CPT investors either trade according to their private information or abstain from trading. Thus, as we move from less to more extreme public beliefs, CPT investors' strategies change. The transitions from one strategy to another can be smooth in the sense that the investor gradually adjusts the probability with which she mixes between two actions, or abrupt, changing discontinuously. The public beliefs at which behavior transitions depend upon an investor's private signal, so that an equilibrium is characterized by four transitions as illustrated in Figure 1. For the case of $\delta > \alpha$, as the public belief increases, the CPT investor with a favorable signal is the first to transition from selling to abstaining. At this price, which I denote p^0 , the equilibrium abruptly changes from a pooling equilibrium in which no information is revealed by a sale, to a separating equilibrium in which the CPT investor with a favorable signal abstains with probability one.⁵ As the public belief increases further, the CPT investor with a favorable signal then smoothly transitions from abstaining to buying over the price range, denoted $p^1 \equiv (\underline{p}^1, \bar{p}^1)$, gradually increasing the probability with which she buys. Within this transition region, CPT investor

⁵The zero profit condition holds at two bid prices for prices just below p^0 . In the equilibrium case of a higher bid price, the CPT investor with a favorable signal sells with probability one. In the other case, she abstains with probability one, lowering the bid price. But, in this latter case, there exists a higher bid price at which the market maker can make strictly positive profits, selling to the CPT investor with the unfavorable signal. The assumption of perfect competition among market makers rules out this possibility.

is characterized by a unique mixing probability because the more the investor buys, the more information is revealed, which in turn raises the ask price reducing the incentive to buy. For the case of $\delta < \alpha$, p^0 and p^1 are defined similarly, but the roles of investors with favorable and unfavorable signals are reversed.

Figure 1: Prospect Theory Investor Behavior in Equilibrium



Note: The upper two plots correspond to $\delta > \alpha$, and the bottom two plots to $\delta < \alpha$. The left two plots illustrate a parameterization for which investors do not trade with either signal over some intermediate range of public beliefs. The right two plots illustrate a second parameterization in which investors instead trade according to private information over this range.

In Figure 1, the upper two plots correspond to $\delta > \alpha$ and the lower two to $\delta < \alpha$. Within each of these two cases, I illustrate the two possible relationships between the locations of the transitions. For the plots on the left of Figure 1, the parameters are such that the two transitions lie on opposite sides of $p = \frac{1}{2}$. In this case, neither type of investor trades over some range of intermediate beliefs. The plots on the right of Figure 1 illustrate a second parameterization in which the transitions lie on the same side of $p = \frac{1}{2}$. In this case, for intermediate public beliefs, a separating equilibrium exists in which CPT investors' trades reveal their private information. Theorem 1 formalizes the illustration of Figure 1, describing the essentially unique equilibrium (up to indifference at transition prices).

Theorem 1: *In equilibrium:*

1. *The market maker posts bid and ask prices (given by (3)) where the conditional buy and sell probabilities are determined by the equilibrium strategies of informed investors that follow.*
2. *For all $p_t \in (0, 1)$, risk-neutral investors buy with favorable signals and sell with unfavorable signals.*
3. *CPT investors' strategies are as follows:*
 - (a) *If $\delta = \alpha$, there exist two cutoff values of loss aversion, $\bar{\lambda} > \underline{\lambda} > 1$ such that, at all $p_t \in (0, 1)$, if $\lambda \leq \underline{\lambda}$, CPT investors buy with favorable signals and sell with unfavorable signals, and, if $\lambda \geq \bar{\lambda}$, they do not trade with either signal. If $\lambda \in (\underline{\lambda}, \bar{\lambda})$, CPT investors mix between buying and not trading with favorable signals and between selling and not trading with unfavorable signals.*
 - (b) *If $\delta \neq \alpha$, for all $\alpha \in (0, 1]$ and $\delta \in (0, 1]$, there exist four transitions, p^0 , $p^1 \equiv (\underline{p}^1, \bar{p}^1)$, and their symmetric counterparts, that characterize strategies:*
 - i. *If $\delta > \alpha$, CPT investors with favorable signals sell for $p_t \leq p^0$, don't trade for $p_t \in (p^0, \underline{p}^1)$, and buy for $p_t \geq \bar{p}^1$. CPT investors with unfavorable signals sell for $p_t \leq 1 - \bar{p}^1$, don't trade for $p_t \in (1 - \underline{p}^1, 1 - p^0)$, and buy for $p_t \geq 1 - p^0$.*
 - ii. *If $\delta < \alpha$, CPT investors with unfavorable signals buy for $p_t \leq p^0$, don't trade for $p_t \in (p^0, \underline{p}^1)$, and sell for $p_t \geq \bar{p}^1$. CPT investors with favorable signals buy for $p_t \leq 1 - \bar{p}^1$, don't trade for $p_t \in (1 - \underline{p}^1, 1 - p^0)$, and sell for $p_t \geq 1 - p^0$.*
 - iii. *Within each of the transitions, p^1 and $1 - p^1$, CPT investors mix between abstaining and trading.*
 - iv. *Transitions do not overlap: $p^0 < \underline{p}^1$ and $\bar{p}^1 < 1 - p^0$, implying $p^0 < \frac{1}{2}$.*

2.6 Analysis: Equilibrium Properties

By the original definition (Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)), buying independently of one's private signal qualifies as herding. However, by more recent definitions of herding and contrarianism in financial markets (Avery and Zemsky (1998) and Park and Sabourian (2011)), CPT investor behavior in the model would not necessarily qualify as herding or contrarianism. In these definitions, herding and contrarian behaviors only arise when investors infer sufficient information from past trades that overwhelms their private information. For example, an investor is said to 'herd' if she would sell at some

initial price (based upon her private information), but instead buys after inferring positive information from others' purchases. A CPT investor that faces a high price instead exhibits behavior that is independent of the past history. If she buys at that price regardless of her signal, she does so whether the price rose (after a series of purchases) or fell from an even higher price (after a series of sales). Of course, high prices are more likely to occur after a sequence of purchases which is why buying will tend to look like trend-following.

In the spirit of the original definition, I define behavior as herding or contrarian if an investor trades independently of her signal. Formal definitions are give in Definition 1 (where I also define the term 'unresponsive' to encompass both).

Definition 1:

1. *An informed investor exhibits herding behavior at the public belief, p_t , if independently of her signal, (i) $p_t > \frac{1}{2}$ and she buys, or (ii) $p_t < \frac{1}{2}$ and she sells.*
2. *An informed investor exhibits contrarian behavior at the public belief, p_t , if independently of her signal, (i) $p_t > \frac{1}{2}$ and she sells, or (ii) $p_t < \frac{1}{2}$ and she buys.*
3. *An informed investor exhibits unresponsive behavior at the public belief, p_t , if her behavior is either herding or contrarian.*

Under these definitions, Figure 1 and Theorem 1 show that CPT investors exhibit herding ($\delta > \alpha$) or contrarian ($\delta < \alpha$) behavior generically (unless $\delta = \alpha$) at sufficiently extreme public beliefs. Rather than stemming from social learning, here herding and contrarian behaviors stem from preferences. When past trends imply highly skewed future returns, CPT investors abandon their private signals and trade in a single direction.

The theory also predicts 'partial' versions of herding and contrarian behaviors (Definition 2) due to the fact that behavior can transition at different public beliefs for different signals. For example, at $p_t > \frac{1}{2}$, an investor may buy with a favorable signal but abstain from trading with an unfavorable signal. Note that, unlike their full version counterparts, these partial versions of herding and contrarian behaviors fully reveal an investor's signal so that informational efficiency is unaffected with a single preference type. When simulating markets, I consider a distribution of CPT preference types. In this case, partial herding or contrarian behaviors may impact informational efficiency because the information revealed by a trade may differ from that revealed by abstention.

Definition 2:

1. *An informed investor exhibits partial herding behavior at the public belief, p_t , if (i) $p_t > \frac{1}{2}$ and she buys with a favorable signal but abstains with an unfavorable signal, or (ii) $p_t < \frac{1}{2}$ and she sells with an unfavorable signal but abstains with a favorable signal.*
2. *An informed investor exhibits partial contrarian behavior at the public belief, p_t , if (i) $p_t > \frac{1}{2}$ and she sells with an unfavorable signal but abstains with a favorable signal, or (ii)*

$p_t < \frac{1}{2}$ and she buys with a favorable signal but abstains with an unfavorable signal.

An immediate consequence of the behavior of risk-neutral investors, provided they arrive with positive probability, is that some information is partially revealed in every period: an *information cascade* in which prices stagnate and subsequent trades reveal no new information never occurs. Thus, by the law of large numbers, public beliefs and bid and ask prices converge to the true asset value in the limit as $T \rightarrow \infty$, achieving full informational efficiency. Conversely, in the absence of risk-neutral investors, information cascades can form in which no information is ever revealed to the market. These cascades can form in one of two ways. First, if the prior is intermediate ($p_1 = \frac{1}{2}$) and the CPT investors are loss averse, they do not trade - a form of non-participation. Second, at extreme prices CPT investors trade unresponsively (herd or act contrarian) so that no information is revealed. In both cases, each CPT investor's decision leaves the price unchanged so that all subsequent investors face the same decision. Their decisions then also reveal no information, and so on. Thus, in a market populated with only CPT investors, even with an infinite sequence of traders, prices never converge to the truth, and may in fact be stuck far from it. With both risk-neutral and CPT investors, loss aversion and unresponsive trades slow, rather than halt, information revelation. I quantify these effects and illustrate an informational cascade in Online Appendix E.

2.7 Analysis: Other Expected Utility Preferences

Having established that CPT can generate both herding and contrarian behaviors, one may wonder whether or not CPT is necessary - can expected utility generate similar behaviors if investors are not risk-neutral? Theorem 2 in Online Appendix A shows that risk-averse expected utility investors neither herd nor act contrarian. Risk-seeking investors, on the other hand, never abstain. The intuition for both results is straightforward. With risk-aversion, herding and contrarianism involve taking a negative expected value bet for one of the private signal realizations, so that abstention is preferable. With risk-seeking, given that a risk-neutral subjects prefers the positive expected value bet of trading with her signal over abstaining, a risk-seeking subject does as well, and so will not abstain. To the extent that most investors are risk-averse, expected utility is unlikely to be a significant driver of unresponsive behavior. However, given that risk-seeking preferences cannot be ruled out as a source of these behaviors, I consider this possibility when analyzing the experimental data.

3 Experiment

3.1 Motivation

Isolating the role of preferences as a source of herding and contrarian behaviors in naturally-occurring settings would be difficult because one would have to rule out inference-based explanations, as well as primitive explanations such as imitation (or anti-imitation). The majority of controlled laboratory studies suffer from the same problems. Because subjects trade sequentially, social factors (the possibility of imitation and the need to make inference from past trades) come into play. Table 2 of Cipriani and Guarino (2009), for example, reports that, as prices become extreme, approximately 20% of trades are herding trades, and 30-40% are contrarian.⁶ This high frequency of contrarian behavior could be driven by subjects' preferences for positive skewness, their desire to go against the crowd, or subjects best-responding to their beliefs that other subjects are herding. Although it may be possible to disentangle these explanations using a structural model, doing so would require strong assumptions about the beliefs subjects have about other subjects' strategies and preferences.

Fortunately, Bisiere, Decamps, and Lovo (2015) develop a more direct way to assess the role of preferences in these environments, through a version of a 'market' in which all decisions are made independently, eliminating social factors. I follow their approach, departing from the theoretical model in one important way. Rather than having subjects trade sequentially, they instead make independent trading decisions. The historical price path a subject observes before deciding how to trade is generated from a sequence of public signal draws, rather than by a market maker who observes previous trades. As such, subjects do not need to form beliefs about the strategies (and preferences) of previous subjects, nor is there scope for them to simply imitate/anti-imitate their predecessors. Thus, any observed herding or contrarian behavior demonstrates that social factors are not necessary for these behaviors. Of course, the cost of this approach is that it abstracts from any interactions between social-based and preference-based reasons for herding, an issue I return to in Section 4.

3.2 Design

I recruited subjects at the University of California, Santa Barbara, from the undergraduate student population over the month of August, 2016, using ORSEE (Greiner(2004)). They participated in a computerized implementation of the experiment implemented in oTree

⁶Cipriani and Guarino (2005) and Drehmann, Oechssler, and Roider (2005) report lower levels of herding, but have low power to detect such behavior because it can only occur in the relatively infrequent cases in which a subject's private signal contradicts the price trend. Cipriani and Guarino (2009) use the strategy method (as do I) to avoid this censoring issue.

(Chen, Schonger, and Wickens (2016)). I conducted three sessions of each of two treatments (MAIN and NO BELIEFS) for a total of 46 subjects in each. Each session of the experiment began by providing subjects with written instructions (see Online Appendix I) which were then read aloud. Each subject took a short quiz to ensure comprehension of the instructions and then participated in 30 rounds of the main decision task. At the completion of all 30 rounds, each subject filled out a short demographics survey (gender, year of university, age) and was then paid the sum of her payoffs across all rounds (in order to make the results more comparable to past experiments such as Cipriani and Guarino (2005,2009)). Average earnings were \$17.13 for an experiment that typically finished in just over an hour.

Each round in the decision task proceeded as follows:

1. Each subject was endowed with $E = 100$ ECUs, experimental currency units.
2. The asset value, V , was randomly chosen by the computer, independently across rounds. With equal probability, it was worth either 100 ECUs or 0 ECUs.
3. The asset was represented by an urn containing either 7 blue balls and 3 green balls (100 ECUs), or 7 green balls and 3 blue balls (0 ECUs), so that the signal precision, $q = 0.7$.
4. Each subject observed a price path of varying length that was determined by a sequence of public draws: one to five balls were drawn, with replacement, from the urn corresponding to the asset's value. Subjects were explicitly told that each price represents the expected value of the asset (what it is worth on average) conditional on the information revealed by past draws. 15 price paths, and their symmetric, counterparts were chosen such that each subject makes multiple decisions at the same price but with different histories (Online Appendix D provides the complete list of price paths). The order of the 30 price paths was randomized across subjects. A screenshot in the instructions provided to subjects (Online Appendix I) depicts the urns and price paths.
5. Subjects made trading decisions only at the final price, p_T (after all public draws).
6. The computer randomly drew a private signal (ball) for the subject. Subjects made trading decisions (buy, sell or not trade) *before* seeing which ball was drawn, one decision for each possible draw. In the MAIN treatment only, subjects were provided correct Bayesian posteriors for the probabilities that the asset value is worth 100 ECUs conditional on *both* the public signal draws and each of their possible private signal draws.
7. Each subject received feedback about the asset value and the realization of her private signal. She received the payoff for the trade corresponding to her realized private signal ($E + V - p_T$ for a buy, $E - V + p_T$ for a sale, and E for no trade).

The only difference between the two treatments is that I provide correct Bayesian posteriors in the MAIN treatment only. Importantly, the framing is identical across treatments, unlike Bisiere, Decamps, and Lovo (2015) who also study the effect of Bayesian posteriors,

but change the framing from a market to a lottery when providing correct posteriors. Thus, the MAIN treatment allows me to cleanly estimate subjects' preferences from their decisions, and the difference across treatments allows me to assess the extent to which Bayesian errors contribute to herding and contrarian behaviors (Bayesian errors being another potential explanation for these behaviors in past experiments). Second, although irrelevant if the theoretical model is correct, I chose to provide subjects with historical price paths. Given potential framing effects (Kahneman and Tversky (1981)), this choice may be important. I purposely chose this framing to make my results more comparable to those of previous market experiments and to real-life investing decisions. In addition, it allows me to explicitly test for any effect of price paths not predicted by the theory. Third, I used the strategy method so that unresponsive behavior can be observed directly (as in Cipriani and Guarino (2009)). Lastly, I follow the previous experimental literature (Cipriani and Guarino (2005) and Drehmann, Oechssler, and Roider (2005)) in having subjects trade at a single price rather than incorporating the bid-ask spread of a fictitious market maker who attempts to infer a subject's private signal from her trade. This procedure both simplifies the problem for subjects, and, more importantly, strengthens incentives to trade according to private information by making the difference between private beliefs and prices larger. Importantly, the predictions of the model are unchanged when trades occur at a single price, except that all transitions in behavior occur at degenerate threshold prices, guaranteeing unique predictions (see Online Appendix B for details).

3.3 Experimental Results

In Section 3.3.1, I present the raw experimental data to assess the extent to which preferences and Bayesian errors are responsible for herding and contrarianism. I also show that particular patterns predicted by the theory are present in the data. In Section 3.3.2, I show that CPT preferences provide a better fit to the data than expected utility, and present estimates of the distribution of preference parameters in the experimental population. In all of the analysis, because subjects make decisions that are independent of each other, the analysis is conducted at the subject level with 46 independent observations per treatment.

3.3.1 The Role of Preferences

In the data analysis, I decided *ex ante* to drop the first 3 markets during which subjects are becoming familiar with the interface and environment. Including this data does not affect any of the qualitative results, nor does restricting the analysis to only the second half of the data. Learning seems to play a very limited role: the results of classifying individual

subjects in Section 3.3.2 are remarkably similar when using only the first or second halves of the data.

I begin by categorizing trading decisions in Table 1. In addition to the definitions from Definition 1 and Definition 2, I define behavior as ‘Revealing’ if a subject buys with a favorable and sells with an unfavorable signal, thus fully revealing her signal. I define ‘Abstain’ as not trading with either signal. Lastly, I define all other trading combinations as ‘Irrational’ given that they can’t be reconciled with the theory (or any other theory I’m aware of). This behavior consists of trading against both signals, always buying or always selling at $p_t = 0.5$, buying with an unfavorable signal but abstaining with a favorable signal, or selling with a favorable signal but abstaining with an unfavorable signal.

Table 1: Behavior By Treatment and Price

Panel A: Aggregated								
Treatment		Abstain	Revealing	Herding	Contrarian	Partial Herding	Partial Contrarian	Irrational
MAIN		5.5	27.5	34.8	10.6	10.8	3.8	7.0
NO		5.2	37.4	13.9	12.8	9.6	16.1	5.2
BELIEFS								
Panel B: By Price								
Treatment	Price	Abstain	Revealing	Herding	Contrarian	Partial Herding	Partial Contrarian	Irrational
MAIN	0.50	2.2	61.6	0.0	0.0	0.0	0.0	36.3
	0.70	2.2	39.7	18.5	10.3	20.9	6.0	2.5
	0.84	5.7	19.6	44.9	12.5	9.8	3.5	4.1
	0.93	8.0	10.1	55.4	12.7	6.2	3.6	4.0
	0.97	15.2	12.0	50.0	14.1	4.4	2.2	2.2
NO	0.50	12.3	63.0	0.0	0.0	0.0	0.0	24.6
BELIEFS								
	0.70	4.9	43.8	10.0	13.0	12.8	12.5	3.0
	0.84	3.8	32.3	13.9	14.4	13.6	19.3	2.7
	0.93	4.0	27.2	23.1	16.3	6.9	20.3	2.2
	0.97	4.4	23.9	21.7	14.1	3.3	29.4	3.3

Note: Percentages of each type of behavior in the MAIN treatment (Bayesian posteriors provided to subjects) and the NO BELIEFS treatment. Panel A contains aggregated statistics. Panel B contains the same information broken down by the normalized price which pools high and low prices together (e.g. 0.7 refers to an actual price of 0.7 or 0.3).

The results of Panel A of Table 1 clearly demonstrate that social factors are not neces-

sary for unresponsive behavior: in each treatment, more than 25% of behavior is herding or contrarian and less than 40% is consistent with the predictions of risk-neutral subjects (revealing trades). To test for treatment differences, I conduct Mann Whitney U (rank-sum) tests that compare the fractions of each type of trade. These tests show that providing subjects with the correct Bayesian posteriors does not reduce unresponsive behavior as we’d expect if Bayesian errors drive this behavior. In fact, the opposite occurs - with correct Bayesian beliefs in the MAIN treatment, we observe significantly *more* herding ($p = 0.000$) with corresponding reductions in revealing ($p = 0.049$) and partial contrarian ($p = 0.000$) trades.⁷ Thus, belief errors appear to work against a preference-based tendency to herd.

Result 1: *Unresponsive behavior is common when social factors are absent, making up 45.4% and 26.7% of trades in the MAIN and NO BELIEFS treatments, respectively. Bayesian errors are not solely responsible for this behavior: providing the correct Bayesian posteriors leads to an increase in herding.*

These results suggest Bayesian errors are not the sole cause of unresponsive behavior, provided subjects actually form beliefs and use them to make trading decisions. If they instead use simple heuristics, we might expect the precise nature of the historical price path to influence behavior, something which would be inconsistent with the theory. To test for this possibility, I ask whether or not the frequencies of herding or contrarianism differ when the price path is (i) monotonic (no contradictory public signals; e.g. three favorable public signals in a row), or (ii) ends in a ‘streak’ (is monotonic or ends in two or more public signals in the same direction; e.g. a sequence of favorable, unfavorable, favorable, favorable).

In performing these tests, I must control for prices because they directly affect decisions, and are mechanically correlated with price paths that are monotonic or have streaks. But, because herding and contrarian behaviors occur above some price, p , and below $1 - p$, they do not depend monotonically on the price. Instead, they theoretically depend on what I call the ‘normalized’ price, $\hat{p}$, a measure of the extremeness of the price: $\hat{p} \equiv p$ if $p \geq \frac{1}{2}$ and $\hat{p} \equiv 1 - p$ if $p < \frac{1}{2}$.⁸ Specifically, I run logit regressions corresponding to the latent variable model, $y_{i,r} = \beta_1 path_r + \beta_2 \hat{p}_r + \varepsilon_{i,r}$, where $y_{i,r}$ is the behavior of interest (herding or contrarian) in round r , $path_r$ is a dummy variable indicating either monotonicity or a streak, and $\hat{p}_r$ represents a set of normalized price fixed effects. I restrict the regression

⁷Under extrapolative expectations (e.g. Barberis, Greenwood, Jin, and Shleifer (2018)), people extrapolate future market returns from past returns, which might lead subjects’ beliefs to be too extreme when not provided with the correct Bayesian posteriors. In this case, we would expect more herding in the NO beliefs treatment, rather than less, as observed. However, note that the fact that the data-generating process is known here may limit the possibility of forming extrapolative expectations.

⁸The tests thus assume subjects behave symmetrically, as in the theory. Because subjects don’t always behave symmetrically (see Section 3.3.2), the results should be interpreted as average effects over high and low prices.

to observations at which unresponsive behavior is possible (normalized prices other than one-half; 1104 observations), clustering standard errors at the subject level.

For herding, I find no statistical evidence of path-dependence in either treatment for either monotonic price paths or streaks ($p > 0.59$ across the four specifications; $p > 0.167$ when partial herding is included). On the other hand, I find some evidence that contrarianism is driven by the nature of the price paths. Monotonic price paths in the MAIN treatment, and streaks in both treatments, significantly *reduce* the frequency of contrarianism ($p < 0.032$), suggesting that subjects find it less desirable to go against the trend in these cases (including partial contrarian trades provides similar results). The sizes of these effects are moderate, however, ranging from a 3.5 to 5% decrease in the probability of a contrarian trade. The evidence is consistent, therefore, with preferences being the primary determinant of behavior.

Result 2: *Herding is not more frequent when price paths are monotonic or end with two or more public signals in the same direction. Contrarian behavior is less likely under both conditions, but only moderately so.*

The theory also predicts particular patterns in unresponsive behavior with respect to prices (Figure 1). Specifically, for almost every preference type (one exception being risk-neutral subjects), revealing trades and abstention should occur at intermediate prices and herding and contrarianism should occur at extreme prices. With a distribution of preference types, we expect these patterns to emerge in the aggregate. To test these predictions, Panel B of Table 1 breaks down the frequency of each type of behavior at each normalized price.

In both treatments, the frequency of revealing behavior almost perfectly monotonically decreases, while the frequency of herding increases, as predicted. I confirm these results via logit regressions of a dummy variable indicating the type of behavior on the normalized price, clustering standard errors at the subject level (the latent variable model is $y_{i,r} = \beta_1 \hat{p}_r + \varepsilon_{i,r}$). For the herding regressions, I exclude a price of one-half where herding is precluded. All four results are highly significant ($p = 0.000$).

Contrarian behavior also increases in each treatment as predicted, but insignificantly so (MAIN, $p = 0.204$; NO BELIEFS, $p = 0.376$), perhaps reflecting the fact that contrarian behavior is not frequent enough to provide sufficient statistical power. Finally, abstention marginally significantly decreases in the NO BELIEFS treatment ($p = 0.047$) as predicted, but increases in the MAIN treatment ($p = 0.005$) because of the tendency of some subjects to abstain at extreme prices. Standard expected utility with a large degree of risk aversion predicts this behavior, and in fact I show in the following section that a handful of subjects are best described by such preferences.


Result 3: *In the aggregate, revealing behavior (always trading according to one's private signal) decreases with the normalized price and herding increases, as predicted by the model.*

Contrarian behavior increases as predicted, but not significantly so. Abstention decreases as predicted in the NO BELIEFS treatment, but increases in the MAIN treatment.

At the individual level, the model predicts that an individual makes herding or contrarian decisions, but not both. To test this prediction, Figure 2 plots the fraction of contrarian decisions versus the fraction of herding decisions for each subject in each treatment. If the theory were perfect, we'd expect 'contrarian' individuals to lie on the y-axis and 'herding' individuals on the x-axis. Although imperfect, most subjects do tend to have a predisposition for one behavior or the other, lying close to one of the axes: across the two treatments 79% of subjects lie within 0.1 of one axis or the other. This predisposition allows me to classify subjects according to their preferences (herding, contrarian, or expected utility) as illustrated in the figure and described in detail in the following section.

Result 4: *The majority of subjects have a predisposition towards either herding or contrarianism, consistent with the model.*

Figure 2: Herding and Contrarian Decisions By Subject



Note: Each point represents a single subject, indicating the fraction of their decisions that are herding or contrarian. Subjects are classified into three types (herding, contrarian, or expected utility) as specified in Section 3.3.2.

3.3.2 Which Preferences Predict Behavior?

The previous section provides evidence consistent with CPT preferences, but does CPT better describe decisions than expected utility? Theorem 2 suggests that expected utility is unlikely to fit the data well because of a limited ability to generate unresponsive trading and abstention, both of which are common in the data. However, it is still possible that the aggregate data could be generated by a mixture of risk-averse and risk-seeking types. To determine whether or not this is the case, I compare the fit of expected utility to that of CPT. I focus on the MAIN treatment where, because Bayesian errors play no role, behavior provides the cleanest expression of preferences. I begin by imposing the same general model on all subjects in the treatment, but allowing the preference parameters to vary on an individual basis.


I consider four candidate preference models. The first is CPT. As I show in Online Appendix B, it has only two degrees of freedom: one relating to loss aversion, λ , and the other relating to the difference between the value function curvature and probability weighting parameters, $\delta - \alpha$. In the model, I constrain α and δ to lie in the interval, $(0, 1]$, which generates restrictions on the trade patterns CPT can generate (see Online Appendix B for details). The second is CPT preferences without loss aversion ($\lambda = 1$), giving it a single degree of freedom and making it directly comparable to the third model, expected utility. For expected utility, in the reported results, I assume CRRA preferences, but the results are virtually identical with CARA preferences. Finally, I consider a non-parametric model which only requires that decisions respect symmetry (the decision at a price p with a favorable signal must match that at a price $1 - p$ with an unfavorable signal). I can construct this ‘best symmetric’ model because of the discrete nature of the prices at which subjects make decisions. It has 27 degrees of freedom and encompasses expected utility, CPT, and any other preference model that respects symmetry as special cases, thus providing an upper bound on the ability of any such model to fit the data (see Online Appendix C for further discussion).

To compare models, I employ the technique of Bisiere, Decamps, and Lovo (2015), calculating a match score for each subject relative to each model of behavior. For a given model and set of preference parameters, I obtain the predicted decisions for each price and for each private signal. I award 0.5 for each of a subject’s 54 decisions that matches the prediction (and zero otherwise), and then divide by the maximum possible score (27) so that the match score lies between 0 and 1.

In practice, because this technique does not allow for probabilistic errors, a range of parameters generates the same prediction for a given model. Thus, in constructing the match score, I first determine the set of predictions consistent with a given model for *any*

preference parameters. For example, in the CPT model, transitions in behavior must occur at two thresholds (and their symmetric counterparts) as illustrated in Figure 1, so it suffices to consider all possible thresholds (see Online Appendix B for details). Once I obtain the thresholds that best match behavior, I then back out the range of parameters consistent with those predictions.

Figure 3: Model Fits in MAIN Treatment



Note: Empirical CDFs of the match scores of subjects for expected utility, prospect theory without loss aversion, prospect theory with loss aversion (thick line), and the best possible model that imposes symmetry.

Figure 3 provides the empirical CDFs of the match scores for each of the four models. To interpret the graph, a y-axis value of 0.5, for example, represents the match score of the median subject for a given model. Thus, the CRRA model matches about 60% of a subject's decisions at the median. The best model that imposes symmetry does better, matching about 75% of a subject's decisions at the median. The fact that it does not match 100% of decisions implies that subjects do not always behave symmetrically at high and low prices (buying is more frequent in both treatments consistent with the findings of Bisiere, Decamps, and Lovo (2015)).

To compare the CDFs, I conduct Kolmogorov-Smirnov(KS) tests where each observation is the match score of an individual subject (46 observations). A particularly striking result

is that CPT explains choices as well as the best possible model that imposes symmetry, even though it has an order of magnitude less degrees of freedom ($p = 0.269$). Furthermore, even restricting CPT to not include loss aversion does almost as well as the full CPT model, and does significantly better than expected utility, both of which have a single degree of freedom ($p = 0.008$).

The previous exercise imposed a single class of preferences on all subjects. To classify individual subjects as specific types (CRRA or CPT), I now allow the class to vary by subject by finding the parameters which provide the best fit (give the highest match score) for each type of preferences. For CPT preferences, I include loss aversion, but the results are identical for all but one subject if loss aversion is excluded.

Taking the preference type that provides a better fit (breaking ties in favor of expected utility), I find that the vast majority of subjects (36/46=78%) are better described by CPT, which is interestingly similar to the 80% of subjects that Bruhin, Fehr-Duda, and Epper (2010) report. Within the set of subjects for which CPT is the best match, 32 subjects (almost 90%) exhibit herding behavior ($\alpha < \delta$). Importantly, however, 4 contrarian ($\alpha > \delta$) subjects exist, which is consistent with CPT but inconsistent with regret theory, the other non-expected utility theory that has been studied in this environment (Qin (2015)).

Among the remaining 10 subjects for which expected utility is a better match, 7 are classified as risk-neutral ($\alpha = \delta = 1$), but 3 subjects are classified as risk-averse (CRRA with $\alpha < 1$), exhibiting behavior that is inconsistent with CPT. These subjects abstain as prices approach zero or one, consistent with a high level of standard risk-aversion. Returning to the graph for the MAIN treatments in Figure 2, we see that subjects classified as herding or contrarian types have a clear disposition for their preferred behavior, with many never making a single decision that contradicts their type.

Result 5: *More than three-quarters of subjects in the MAIN treatment are better described by prospect theory than expected utility (CRRA) and, for these, herding strategies are dominant. Furthermore, prospect theory fits the data as well as any model that imposes symmetry.*

Because δ and α do not affect behavior independently, I can associate a range of CPT parameters with each of the 36 individuals best fit by CPT preferences. In order to conduct simulation exercises using the distribution of CPT preferences from the experiment, I choose the parameters closest to the risk-neutral benchmark ($\delta = \alpha = 1$), normalizing $\delta = 1$ for herding types and $\alpha = 1$ for contrarian types. For the 32 herding types, the largest possible α consistent with behavior ranges from 0.25 to 0.75. For the 4 contrarian types, the largest possible δ ranges from 0.5 to 0.66. For 34/36 of these subjects, the loss-aversion parameter, λ , lies between 1 and 2, consistent with a large literature.

The median and modal subject has $\delta = 1$ and $\alpha = 0.5$, corresponding to herding at prices, $p \geq 0.84$ and $p \leq 0.16$. The original finding of Kahneman and Tversky (1992) that $\delta < \alpha$ has led to a belief that probability weighting dominates value function curvature, in contrast to my findings. However, a closer examination of the literature shows that parameter estimates vary widely and my finding is nothing out of the ordinary: across the ten preference measurement studies summarized in Table A.3 of Glimcher and Fehr (2013) that use the same probability weighting function I use, four also find $\delta > \alpha$ at the median. Unfortunately for comparison purposes, very little evidence on the overall joint distribution exists. Only a single paper I’m aware of (Zeisberger, Vrecko, and Langer (2012)) reports individual preference parameters. They document substantial heterogeneity as I do, but, importantly, also that estimates for the majority of subjects are stable across elicitation tasks one month apart.

In Online Appendix E, I plot the joint distribution of the preference parameters and then simulate a market composed of individuals with preferences drawn from it. I document substantial slowdowns in the rate at which prices converge to the true asset value, relative to a market composed of only risk-neutral individuals, due to the fact that herding and contrarian behaviors do not reveal information to the market. In Online Appendix F, I use the distribution of preferences as well as the differences in behavior across the MAIN and NO BELIEFS treatments to assess the types of Bayesian errors make in the NO BELIEFS treatment. I find that subjects underinfer the information in public signals, consistent with the conclusion of Benjamin (2018) that, across many studies, underinference from information is the norm.

4 Discussion

Herding and contrarian strategies are of interest because of their potential for reducing the informational efficiency of prices by not revealing information to the market. These strategies are often thought to rely on (or even be defined as relying on) the observation of past histories of actions and/or prices. In this paper, however, I have shown both theoretically and via a laboratory experiment that they can emerge from preferences. The strategies depend only upon past price trends to the extent that these trends are related to *future* returns. In the experiment, I control for social and belief-based causes of behavior, yet observe economically significant levels of herding and contrarianism, demonstrating a clear role for preferences. Among the three-quarters of subjects whose behavior is more consistent with prospect theory

than expected utility, almost 90% of subjects exhibit a preference for herding strategies.⁹

This high proportion of herding is consistent with a long-standing narrative of financial markets - investors chase price trends, producing bubbles and crashes (Mackay (1841)). In the context of previous experimental findings, however, the results are surprising. Cipriani and Guarino (2005,2009) and Drehmann, Oechssler, and Roider (2005) each report more contrarian than herding behavior. How to reconcile these findings and to what extent may prospect theory preferences be responsible for behavior in past experiments? To answer these questions, consider the experiment of Cipriani and Guarino (2009), which is notable because they conduct their experiment with a sample of financial market professionals. Thus, the extent to which my data corresponds to theirs provides some assurance that my results are not particular to an undergraduate student population. Table 2 compares the frequencies of herding and contrarian decisions in my NO BELIEFS treatment (subjects in Cipriani and Guarino (2009) had to perform Bayesian updating) with those reported in Cipriani and Guarino (2009), Table 2.

Table 2: NO BELIEFS versus Cipriani and Guarino (2009)

Normalized Price	Herding		Contrarian	
	NO BELIEFS	CG2009	NO BELIEFS	CG2009
0.50	0.0	0.0	0.0	0.0
0.70	10.0	5.7	13.0	12.9
0.84	13.9	16.1	14.4	26.6
0.93	23.1	23.9	16.3	30.4
≥ 0.97	21.7	21.9	14.1	40.6

Note: Percentages of herding and contrarian decisions in the NO BELIEFS treatment and from Table 2 in Cipriani and Guarino (2009), CG2009.

In Table 2, we see that the frequencies of herding at each price are remarkably similar across the two experiments, and are not insignificant, exceeding 20% of decisions as prices become extreme. This finding suggests that prospect theory may be responsible for a significant fraction of decisions in Cipriani and Guarino (2009) because no other available theory

⁹In Kendall (2020), I also observe frequent herding, but because the environment is more complex it is difficult to draw strong conclusions about the mechanism behind this behavior.

generates herding in this environment. On the other hand, the level of contrarianism is much higher in Cipriani and Guarino (2009), which leads to an intriguing possibility. In their setup, subjects had to form beliefs about the strategies of other subjects. *If* subjects understand that some fraction of subjects herd (an understanding that experienced financial professionals may be more likely to have), it can be a best response to act contrarian because prices are too extreme, on average. Understanding the interaction between preference-based herding and contrarianism and social learning (i.e. whether or not subjects form and best respond to correct beliefs about the preferences of others) in this (and related) environments is an interesting question for future research.

Interpreted as a preference for negative skewness, my results are consistent with those of Huber, Palan, and Zeisberger (2019). They document that positively-skewed assets are more likely to be perceived as risky (having a higher probability of a loss) and thus trade at lower prices than negatively-skewed assets. In real-world markets, the evidence has tended to point to preferences for positive skewness (Kumar (2009), Green and Hwang (2012), Eraker and Ready (2015)). Unlike the laboratory evidence, however, this evidence reflects selection into these markets (as documented by Kumar (2009)).

Prospect theory is, at its core, about framing - people perceive gambles differently when framed as gains or losses (Kahneman and Tversky (1979,1981)). Financial markets frame choices in a particular way, one which I adopted in the experimental design. But, the potential for framing effects raises the possibility that the results of the experiment would be different if framed in some other way. For example, ignoring private information, at a price of 0.9, a subject choosing to buy or sell faces a choice between two lotteries that result in similar payoffs when $V = 1$ (1.1 vs. 0.9), but very different payoffs when $V = 0$ (0.1 vs. 1.9). If framed as a lottery, the state $V = 1$ may be salient. If so, a subject may sell because 1.9 is preferable to 0.1, producing contrarian behavior. Formalizing this idea, one can replace the CPT investors of the model with investors that satisfy the salience model of Bordalo, Gennaioli, and Shleifer (2012). In this case, investors will sometimes follow contrarian strategies, but never herd. Consistent with this conjecture, Bisiere, Decamps, and Lovo (2015) find more herding when decisions are framed as a market in their robustness treatment (SME) than as a choice between lotteries (LE). If framing matters, the way in which market data is presented may affect the extent of herding and thus the informational efficiency of the market. Exploring this possibility is an interesting, and potentially very important, topic for future research.

A broader theme suggested by this paper is that prospect theory preferences can play a critical role in market and game experiments. In interpreting the data from these types of experiments, risk-neutral (expected utility) preferences are often assumed, citing the Ra-

bin critique (Rabin (2000)): *if* people are expected utility maximizers, then they must be risk-neutral in the laboratory. But, Rabin’s conclusion is not that we should assume risk-neutrality. Instead, he argues that because we often see departures from risk-neutrality over small stakes, we should consider models other than expected utility. Failing to do so can lead to very erroneous conclusions about the underlying reasons for any observed departures from ‘rational’ behavior.

References

- [1] Avery, Christopher and Peter Zemsky (1998), “Multidimensional uncertainty and herd behavior in financial markets”, *American Economic Review*, 88 (4), 724-748
- [2] Baltzer, Markus, Stephan Jank, and Esad Smajlbegovic (2019), “Who trades on momentum?”, *Journal of Financial Markets*, 42, 56-74
- [3] Banerjee, Abhijit (1992), “A simple model of herd behavior”, *The Quarterly Journal of Economics*, 107 (3), 797-817
- [4] Baillon, Aurélien, Han Bleichrodt, and Vitalie Spinu (2019), “Searching for the reference point”, *Management Science*, 66 (1), 93-112
- [5] Barberis, Nicholas (2012), “A model of casino gambling”, *Management Science*, 58 (1), 35-51
- [6] Barberis, Nicholas (2013), “Thirty years of prospect theory in economics: a review and assessment”, *Journal of Economic Perspectives*, 27 (1), 173-196
- [7] Barberis, Nicholas and Ming Huang (2008), “Stocks as lotteries: the implications of probability weighting for security prices”, *American Economic Review*, 98 (5), 2066-2100
- [8] Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer (2018), “Extrapolation and bubbles”, *Journal of Financial Economics*, 129, 203-227
- [9] Benjamin, Daniel J. (2018), “Errors in probabilistic reasoning and judgment biases,” In *Handbook of Behavioral Economics* (edited by Douglas Bernheim, Stefano DellaVigna, and David Laibson), Elsevier Press
- [10] Bisière, Christophe, Jean-Paul Décamps, and Stefano Lovo (2015), “Risk attitude, beliefs updating, and the information content of trades: an experiment”, *Management Science*, 61 (6), 1378-1397

- [11] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), “A theory of fads, fashion, custom, and cultural change as informational cascades”, *The Journal of Political Economy*, 100 (5), 992-1026
- [12] Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2012), “Salience theory of choice under risk”, *The Quarterly Journal of Economics*, 127 (3), 1243-1285
- [13] Boortz, Christopher (2016), “Irrational Exuberance and Herding in Financial Markets”, SFB 649 Discussion Paper 2016-016
- [14] Bond, Philip, Alex Edmans, and Itay Goldstein (2012), “The real effects of financial markets”, *Annual Review of Financial Economics*, 4, 339-360
- [15] Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010), “Risk and rationality: uncovering heterogeneity in probability distortion”, *Econometrica*, 78 (4), 1375-1412
- [16] Chen, Joseph, Harrison Hong, and Jeremy C. Stein (2001), “Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices”, *Journal of Financial Economics*, 61, 345-381
- [17] Chen, Daniel L., Martin Schonger, and Chris Wickens (2016), “oTree - An Open-Source Platform for Laboratory, Online and Field Experiments”, *Journal of Behavioral and Experimental Finance*, 9, 88-97
- [18] Cipriani, Marco and Antonio Guarino (2005), “Herd behavior in a laboratory financial market”, *American Economic Review*, 95 (5), 1427-1443
- [19] Cipriani, Marco and Antonio Guarino (2008), “Transaction costs and informational cascades in financial markets”, *Journal of Economic Behavior & Organization*, 68, 581-592
- [20] Cipriani, Marco and Antonio Guarino (2008), “Herd behavior and contagion in financial markets”, *The B.E. Journal of Theoretical Economics*, 8 (1), Art.24
- [21] Cipriani, Marco and Antonio Guarino (2009), “Herd behavior in financial markets: an experiment with financial market professionals”, *Journal of the European Economic Association*, 7 (1), 206-233
- [22] De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann (1990), “Positive feedback investment strategies and destabilizing rational speculation”, *The Journal of Finance*, 45 (2), 379-395

- [23] Décamps, Jean-Paul and Stefano Lovo (2008), “Informational cascades with endogenous prices: the role of risk aversion”, *Journal of Mathematical Economics*, 42, 109-120
- [24] Devenow, Andrea and Ivo Welch (1996), “Rational herding in financial economics”, *European Economic Review*, 40, 603-615
- [25] Drehmann, Mathias, Jörg Oechssler, and Andreas Roider (2005), “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment”, *American Economic Review*, 95 (5), 1403-1426
- [26] Ford, J. L., David Kelsey, and Wei Pang (2013), “Information and ambiguity: herd and contrarian behaviour in financial markets”, *Theory and Decision*, 75 (1), 1-15
- [27] Eraker, Bjørn and Mark Ready (2015), “Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks”, *Journal of Financial Economics*, 115 (3), 486-504
- [28] Glimcher, Paul and Ernst Fehr (2013), “Neuroeconomics: decision making the brain”, Academic Press, 2nd edition
- [29] Glosten, Lawrence, R. and Paul R. Milgrom (1985), “Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders”, *Journal of Financial Economics*, 14, 71-100
- [30] Green, T. Clifton and Byoung-Hyoun Hwang (2012), “Initial public offerings as lotteries: skewness preference and first-day returns”, *Management Science*, 58 (2), 432-444
- [31] Robin Greenwood and Andrei Shleifer (2014), “Expectations of returns and expected returns”, *Review of Financial Studies*, 27, 714-746
- [32] Greiner, Ben (2004), “The online recruitment system ORSEE - a guide for the organization of experiments in economics”, Papers on Strategic Interaction 2003-10, Max Planck Institute of Economics, Strategic Interaction Group
- [33] Grinblatt, Mark, Gergana Jostova, Lubomir Petrasek, and Alexander Philipov (2016), “Style and skill: hedge funds, mutual funds, and momentum”, mimeo, SSRN, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2712050
- [34] Grinblatt, Mark, Sheridan Titman, and Russ Wermers (1995), “Momentum investment strategies, portfolio performance, and herding: a study of mutual fund behavior”, *American Economic Review*, 85 (5), 1088-1105

- [35] Grinblatt, Mark, and Matti Keloharju (2000), “The investment behavior and performance of various investor types: a study of Finland’s unique data set”, *Journal of Financial Economics*, 55 (1), 43-67
- [36] Harvey, Campbell R. and Akhtar Siddique (2000), “Conditional skewness in asset pricing tests”, *The Journal of Finance*, LV (3), 1263-1296
- [37] Hirshleifer, David and Siew Hong Teoh (2003), “Herd behavior and cascading in capital markets: a review and synthesis”, *European Financial Management*, 9 (1), 25-66
- [38] Huber, Juergen, Stefan Palan, and Stefan Zeisberger (2019), “Does investor risk perception drive asset prices in markets? Experimental evidence”, *Journal of Banking and Finance*, 108, 105635
- [39] Kahneman, Daniel and Amos Tversky (1981), “Prospect theory: an analysis of decision under risk”, *Econometrica*, 47 (2), 263-291
- [40] Kahneman, Daniel and Amos Tversky (1981), “The framing of decisions and the psychology of choice”, *Science*, 211 (4481), 453-458
- [41] Kahneman, Daniel and Amos Tversky (1992), “Advances in prospect theory: cumulative representation of uncertainty”, *Journal of Risk and Uncertainty*, 5, 297-323
- [42] Kendall, Chad (2020), “Market panics, frenzies, and informational efficiency: theory and experiment”, *American Economic Journal: Microeconomics*, 12 (3), 76-115
- [43] Kőszegi, Botond and Matthew Rabin (2006), “A model of reference-dependent preferences”, *The Quarterly Journal of Economics*, CXXI (4), 1133-1165
- [44] Kőszegi, Botond and Matthew Rabin (2007), “Reference-dependent risk attitudes”, *American Economic Review*, 97 (4), 1047-1073
- [45] Kumar, Alok (2009), “Who gambles in the stock market?”, *The Journal of Finance*, 64, 4, 1889-1933
- [46] Mackay, Charles (1841), “Memoirs of extraordinary popular delusions and the madness of crowds”, (1 ed.). London: Richard Bentley
- [47] Park, Andreas and Daniel SgROI (2012), “Herding, contrarianism, and delay in financial market trading”, *European Economic Review*, 56, 1020-1037
- [48] Park, Andreas and Hamid Sabourian (2011), “Herding and contrarian behavior in financial markets”, *Econometrica*, 69 (4), 973-1026

- [49] Qin, Jie (2015), “A model of regret, investor behavior, and market turbulence”, *Journal of Economic Theory*, 160, 150-174
- [50] Rabin, Matthew (2000), “Risk-aversion and expected-utility theory: a calibration theorem”, *Econometrica*, 68 (5), 1281-1292
- [51] Roth, Alvin E. and Michael W. Malouf (1979), “Game theoretic models and the role of information in bargaining”, *Psychological Review*, 86, 575-594
- [52] Zeisberger, Stefan, Dennis Vrecko, and Thomas Langer (2012), “Measuring the time stability of prospect theory preferences”, *Theory and Decision*, 72, 359-386