

Does Belief Heterogeneity Explain Asset Prices: The Case of the Longshot Bias

Amit Gandhi and Ricardo Serrano-Padial*

September 17, 2012

Abstract

This paper studies belief heterogeneity in a benchmark competitive asset market: a market for Arrow-Debreu securities. We show that differences in agents' beliefs lead to a systematic pricing pattern, the favorite longshot bias (FLB): securities with a low payout probability are overpriced while securities with high probability payout are underpriced. We apply demand estimation techniques to betting market data, and find that the observed FLB is explained by a two-type population consisting of *canonical* traders, who hold virtually correct beliefs and are the majority type in the population (70%); and *noise* traders exhibiting significant belief dispersion. Furthermore, using formal model comparisons and also exploiting variation in public information across markets in our dataset, we show that our belief heterogeneity model empirically outperforms existing preference-based explanations of the FLB, such as risk-loving or prospect theory.

JEL Classification: *C13, C51, D40, G13, L00.*

Keywords: *heterogeneity, prospect theory, favorite-longshot, rational expectations, demand estimation, random utility, noise traders, risk preferences.*

1 Introduction

There are two important issues that have gained prominence in the study of financial markets. First, there is a growing consensus that the sheer trading volume we observe cannot be explained without resorting to information-driven trade (Cochrane, 2007), i.e., belief heterogeneity is likely a major source of gains from trade in financial markets. Second, there is also a growing body of evidence that documents “anomalies” in prices (see e.g., Keim, 2008), i.e., systematic patterns where prices do not accurately reflect the underlying fundamentals of securities.

*Department of Economics, University of Wisconsin-Madison. Contact: agandhi@ssc.wisc.edu and rserrano@ssc.wisc.edu. This paper was earlier circulated under the title “The Aggregate Effects of Idiosyncratic Heterogeneity: Evidence from an Arrow-Debreu Securities Market.” We are grateful to Enrique Sentana and three anonymous referees, whose comments greatly benefited the paper, and to Ken Hendricks, Jean-Francois Houde, Marco Ottaviani, Antonio Penta, Dan Quint, Marzena Rostek and Marek Weretka for very useful suggestions. We also thank Zhentong Lu for outstanding research assistance.

In this paper we show these two issues are closely related: belief heterogeneity can serve as a natural foundation for pricing patterns that depart from the predictions of standard neoclassical theory. We focus on a particular pricing ‘anomaly’, known as the favorite-longshot bias (FLB), which has been found in a number of market settings, particularly in betting markets.¹ Betting markets are important for several reasons. Chief among them is that they represent real world Arrow-Debreu security markets, which play a foundational role in the theoretical study of financial markets. In addition, unlike in other markets, we observe ex-post security returns which, combined with the availability of large datasets, allows tracing prices to fundamentals.² There is a long empirical literature that widely documents the existence of the FLB in these markets – securities with a large probability of payout (favorites) yield higher average returns than securities with a small payout probability (longshots), with observed disparities in returns as high as 200%.³ Because of this, the FLB is often viewed as evidence against the standard notion of agents being *risk averse, expected utility* maximizers with *correct beliefs* (or rational expectations) about the underlying fundamentals (Thaler and Ziemba, 1988). This has led to two main alternative approaches to explain the observed pattern of returns, both of them based on departures from standard preferences: the first one introduces risk loving (Quandt, 1986); while the second departs from expected utility by assuming rank dependent utility or cumulative prospect theory preferences (Jullien and Salanié, 2000; Snowberg and Wolfers, 2010).

The purpose of this paper is to study, both theoretically and empirically, the relationship between belief heterogeneity and pricing in competitive asset markets. We use the standard Arrow-Debreu (A-D) security market setting and explore whether belief heterogeneity, apart from explaining trade, can also explain the existence of the FLB. In addition, we investigate whether belief heterogeneity empirically outperforms existing preference-based explanations of the FLB, both in terms of explaining the overall variation of prices in betting data, and in terms of predicting the reaction of prices to changes in the availability of information in the market. Our main findings can be summarized as follows:

- We theoretically show that belief heterogeneity among risk neutral traders naturally leads to the FLB in competitive A-D security markets: as long as the correct beliefs about the true payout probabilities lie in the support of the population belief distribution, we should expect favorites to be underpriced and longshots to be overpriced.

¹There is evidence of the FLB in betting and prediction markets (see e.g., Griffith (1949); Sauer (1998); Jullien and Salanié (2000); Snowberg and Wolfers (2010) for a review), option markets (where the FLB is associated to the “volatility smile”), and in some derivative markets (Tompkins et al., 2008).

²An additional key feature of betting markets is the absence of differences in liquidity or transaction cost across securities. This is in contrast with other markets where pricing patterns are closely associated to differences in liquidity and trading costs (see e.g. Sadka and Scherbina (2007) and references therein).

³See, for instance, Jullien and Salanié (2008) for a review of the literature.

- We present a framework for estimating heterogeneous agent models in A-D markets using aggregate data on prices and ex-post returns. In particular, we adapt demand estimation techniques for random utility models (e.g. Berry et al., 1995) to financial markets, which allow us to accommodate heterogeneity in both risk preferences and beliefs. Using betting market data (horse racing data), we estimate a model of (idiosyncratic) belief heterogeneity. We find that a two-type population of risk-neutral traders, in which the prevalent type (about 70% of agents) holds virtually correct beliefs with minimal belief dispersion and the other type exhibits higher belief dispersion, explains the observed pricing pattern remarkably well. This provides suggestive evidence of the existence of two types of agents in these markets, informed and *noise* traders.
- We compare our model to the two leading preference-based approaches to explaining the FLB: risk loving and non-expected utility. In order to do so, we estimate a model with a population of traders exhibiting heterogeneous risk attitudes (which includes risk loving as a special case), and also a representative agent model with preferences that follow cumulative prospect theory (which also nests rank-dependent utility). We find that the application of these alternative theories to our data yields parameter estimates that are similar to existing results in the literature (most notably Jullien and Salanié (2000) and Snowberg and Wolfers, 2010). However, we also find that our heterogeneous beliefs model explains the price variation in the data substantially better than these preference-based approaches.⁴
- In order to further understand the role of belief heterogeneity in the data, we exploit a source of variation in the availability of information at races. Specifically, races that are run at the same track on the same day take two different forms, maiden or non-maiden. Because only experienced horses can participate in non-maiden races and new horses (without any racing history) participate in maiden races, traders in non-maiden races have access to richer information about the underlying value of the A-D securities associated to each horse. We find that this difference in the information structure across race types is reflected very distinctly in the asset prices: maiden races exhibit a much more pronounced FLB as compared to non-maiden races. We estimate our heterogeneous beliefs model on each of these subsamples of races and show that the model explains this pricing difference in a natural way: the same proportions of informed and noise traders are present in both types of races, but

⁴We also find support for our risk neutral assumption: we estimate a belief heterogeneity model with CARA agents and find that the CARA coefficient is not significantly different from zero. This is consistent with agents being essentially risk neutral when stakes are low (Bombardini and Trebbi, 2010).

traders simply have more dispersed beliefs in the low information races. Preference based theories, on the other hand, cannot easily account for this change in the pattern of returns across information environments.

What is the intuition for why belief heterogeneity empirically explains the FLB so readily? Belief heterogeneity can generate a dramatic disparity in returns even when beliefs are heavily concentrated around the correct beliefs because the FLB, like many pricing puzzles observed in financial markets, is fundamentally driven by behavior at the tails of the underlying value distribution. Most of the mispricing in the data is associated to extreme longshots: disparity in expected returns becomes most apparent for securities with payout probabilities less than 1%. Accordingly, a little dispersion in beliefs suffices to induce enough demand on those longshots to generate substantial overpricing in equilibrium. For example, in our data, the average security in the data has over twice as high an expected return as compared to extreme longshots that payout with less than 1 percent probability. In our estimated model, only roughly 5 percent of agents would be willing to invest in these longshot assets relative to the average security. But this small demand is sufficient to generate the degree of observed overpricing found in the data.

Finally, our findings highlight the potentially misleading inferences that can be drawn from representative agent models. Although a representative agent that rationalizes the aggregate demand of a population of heterogeneous agents may exist, this representative agent's preferences need not reflect the preferences in the population in any meaningful sense.⁵ For example, in our estimated model of belief heterogeneity, the modal behavior is well captured by the textbook homogeneous beliefs and risk neutrality model. In contrast, a representative agent model would require highly non-standard preferences or beliefs to explain the same data.

Related Literature Our theoretical model is related to existing models linking belief heterogeneity to the FLB, notably the competitive model of Ali (1977) and the asymmetric information approach of Ottaviani and Sorensen (2006, 2010b).⁶ Unlike the latter papers, our goal is not to provide an informational foundation to heterogeneous posterior beliefs in asset markets, but rather find general conditions on the distribution of such beliefs that lead to the FLB. In terms of theoretical predictions, the main difference between our approach and the existing literature is our focus on price behavior at the tails of the value

⁵The theoretical analysis of Ottaviani and Sorensen (2010a) also stresses the potential divergence between representative and heterogeneous agent interpretations of prices. In particular, they show that a population of perfectly Bayesian traders with heterogeneous beliefs leads to equilibrium prices that can only be generated by a non-Bayesian representative agent.

⁶See also the work of Shin (1991, 1992) and Potters and Wit (1996).

distribution, which is where the FLB empirically arises. In contrast, existing research typically defines the FLB for the whole range of values, e.g., by predicting returns to be strictly increasing in payout probabilities.⁷ Accordingly, we are able to substantially weaken the conditions under which the FLB arises.

Our estimation strategy builds on the work of Jullien and Salanié (2000), who propose a maximum likelihood framework to estimate representative agent preferences using data on prices and ex-post returns. Our insight is that we can add heterogeneity to this framework by combining it with the approach of Berry et al. (1995), originally developed to estimate demand for differentiated products from aggregate data on prices and product market shares. The particular random utility model we propose for belief heterogeneity is also new, however, to the demand estimation literature. Finally, this paper adds a novel dimension to the empirical literature on the FLB, which, to date, has emphasized preference-based explanations of the observed pricing patterns (Jullien and Salanié, 2000; Snowberg and Wolfers, 2010; Golec and Tamarkin, 1998). Belief based trade has not yet been explored as an empirical alternative, which is an important gap we fill. Relatedly, our finding that the magnitude of the FLB changes across races characterized by different information structures, which is a prediction of our belief heterogeneity model, appears to be a new contribution to the empirical literature on betting and financial markets.

The plan of the paper is as follows. In Section 2, we describe the market and characterize the demand of a risk neutral agent. In Section 3, we introduce belief heterogeneity and show that it produces the FLB in equilibrium. In Section 4, we describe the dataset and illustrate the observed price pattern. We show in Section 5 how to move from the theory to an empirical framework for measuring heterogeneity. In Section 6, we describe our estimation results, while we compare our approach to the leading preference-based alternatives in Section 7. Section 8 concludes.

2 The Market

We consider a competitive market for n Arrow-Debreu securities. Security i pays \$1 if outcome i takes place. Before the market opens, nature determines the state $\mathbf{p} = (p_1, \dots, p_n) \in \text{int}\Delta^{n-1}$, where $\text{int}\Delta^{n-1}$ is the interior of the $(n - 1)$ dimensional simplex and p_i is the probability that outcome i is realized. That is, Δ^{n-1} is the set of all possible probability distributions and the state $\mathbf{p} \gg 0$ is a particular probability distribution over the n possible outcomes associated to the A-D securities.

⁷Monotonicity of returns is predicted by Quandt (1986). In a market with two A-D securities Ali (1977) and Ottaviani and Sorensen (2006, 2010b) provide conditions for the security with the higher (lower) probability to be underpriced (overpriced).

Let ρ_i denote the price of security i . We allow for the possibility of positive transaction costs represented by a fraction $\tau \geq 0$ of each dollar invested in the market that the institution keeps for its own profit. In this context, the expected gross return of investing \$1 in security $i = 1, \dots, n$ is given by

$$ER_i = (1 - \tau) \frac{p_i}{\rho_i}, \quad (1)$$

and the expected net return is $ER_i - 1$. Accordingly, securities that are underpriced relative to their true chance of yielding returns have a higher expected return as compared to securities that are overpriced.

The market consists of a population $T = [0, 1]$ with a continuum of risk neutral agents with finite endowments. Assume for the moment that there are no transaction costs ($\tau = 0$) and that the outside option yields zero net returns given agent's beliefs. The demand of agent $t \in T$ is a bundle (x_1, \dots, x_n) such that $\sum x_i \leq w_t$, where $x_i \geq 0$ is the amount invested on security i , and $w_t > 0$ is the dollar endowment of the agent.

Due to risk neutrality, the demand for security i of any agent t with well-defined posterior beliefs, denoted by $(\pi_{1t}, \dots, \pi_{nt})$, is determined by the relative comparisons of *subjective* expected returns, given by π_{it}/ρ_i , $i = 1, \dots, n$. That is, agent t will invest on the security i if it yields the highest subjective return, i.e., if $\pi_{it}/\rho_i > \pi_{jt}/\rho_j$ for all $j \neq i$.

Let $r_{it} = p_i/\pi_{it}$, which reflects agent t 's relative deviation from correct beliefs about security i . The next lemma shows that the demand of agent t for security i can be characterized instead in terms of the *objective* expected returns and a vector of *indifference ratios*, given by r_{it}/r_{jt} for $j \neq i$: the agent is indifferent between security i and j whenever the ratio of expected returns ER_i/ER_j is equal to the indifference ratio r_{it}/r_{jt} .

Lemma 1. *There exist a vector $\mathbf{r}_t = (r_{1t}, \dots, r_{nt}) > 0$ with $r_{it} = p_i/\pi_{it}$ such that, for all $(\rho_1, \dots, \rho_n) \gg 0$ and all $i = 1, \dots, n$, the demand of a risk neutral agent t satisfies*

$$q_{it}(ER_1, \dots, ER_n) = \begin{cases} w_t & \frac{ER_i}{ER_j} > \frac{r_{it}}{r_{jt}} \quad \forall j \neq i \\ \{x : x \leq w_t\} & \frac{ER_i}{ER_j} \geq \frac{r_{it}}{r_{jt}} \quad \forall j, \frac{ER_i}{ER_j} = \frac{r_{it}}{r_{jt}} \text{ some } j \neq i \\ 0 & \frac{ER_i}{ER_j} < \frac{r_{it}}{r_{jt}} \text{ some } j \neq i, \end{cases}$$

and

$$\sum_i q_{it}(ER_1, \dots, ER_n) \leq w_t,$$

which holds with equality if $\frac{ER_i}{ER_j} \neq \frac{r_{it}}{r_{jt}}$ for some i, j .

Proof. In the *Appendix*. □

Notice that an agent with correct beliefs has indifference ratios equal to one for all i, j . Likewise, if an agent has an indifference ratio $r_{1t}/r_{2t} = 2$, then asset 1 would have to pay twice as high objective expected returns as asset 2 for the agent to be indifferent between the two (thus the agent either seriously overestimates the probability of asset two paying off, or underestimates asset 1).

There are two advantages of working with objective returns and indifference ratios rather than dealing directly with prices and subjective beliefs. First, the fact that indifference ratios reflect relative differences with respect to the true underlying payout probabilities allows us to define (idiosyncratic) belief heterogeneity in terms of deviations from correct beliefs (see *Assumption 1* below). Second, the indifference conditions in *Lemma 1* readily lead to the random utility model we use later in our empirical analysis.

Remark. We assume risk neutrality for two reasons. First, unlike risk aversion, it rationalizes the discrete choice of investing all the endowment in one of the securities, rather than hedging or investing only a fraction of it.⁸ Second, empirical studies have shown that people tend to be risk-neutral when stakes are low (e.g. Bombardini and Trebbi, 2010), as it is often the case for most traders in betting markets. Nonetheless, we allow for alternative risk attitudes in our empirical estimation and find support for risk neutrality in the presence of belief heterogeneity (see Section 6).

Finally, let s_i be security i 's market share, i.e., the amount invested on security i relative to the total (finite) amount invested in the market. Market clearing implies that prices equal market shares, as long as agents, when indifferent between investing or not, choose to invest either all their endowment or nothing.⁹ It is worth noting that, by definition, prices are equal to shares in parimutuel markets, as are the markets comprising our dataset.

3 Belief Heterogeneity and Prices

In a standard representative agent model, traders are homogeneous in terms of beliefs and preferences. In particular, they hold correct posterior beliefs about the state of the world. We call such (risk-neutral) agents *canonical traders*, and their individual demand

⁸Since endowments are not observed, most empirical studies of the FLB take a discrete choice approach. See for instance Jullien and Salanié (2000); Snowberg and Wolfers (2010).

⁹Market clearing means the supply of dollars equals the demand of dollars in each of the possible n outcomes. This happens if and only if

$$\frac{q_i}{\rho_i} = (q_1 + \dots + q_n) \iff \rho_i = s_i = \frac{q_i}{q_1 + \dots + q_n} \quad (\forall i),$$

where q_i represents the aggregate investment in dollars on security i . If a mass of agents is indifferent between investing in some securities and staying out of the market, and they invest only a fraction of their endowment then the supply of dollars may be higher than the total investment in the market, in which case some agents are short-selling some of the securities.

is characterized by $\mathbf{r}_t = (1, \dots, 1)$. In this context, the only trading equilibrium outcome is $ER_i = ER_j$ for all i, j , since otherwise the market does not clear.¹⁰ That is, there do not exist any gains from trade in this economy and agents are indifferent between investing in any of the securities or staying out of the market and thus such equilibrium is not robust to the introduction of transaction costs.

Now consider introducing gains from trade in this baseline model by letting agents exhibit differences in posterior beliefs, i.e., heterogeneity in indifference ratios. Let the mass function of \mathbf{r}_t be given by the conditional probability measure $P[\cdot \mid \mathbf{p}, \theta]$, which depends both on the state of the world p and possibly on other characteristics of the market $\theta \in \Theta$. We only require that there is always a positive mass to “both” sides of the canonical trader —this avoids the no-trade trap discussed above. To capture this condition, given $z \in \mathbb{R}_+^{n-1}$, let

$$L_i[z \mid \mathbf{p}, \theta] = P \left[\left(\frac{r_{it}}{r_{1t}}, \dots, \frac{r_{it}}{r_{i-1t}}, \frac{r_{it}}{r_{i+1t}}, \dots, \frac{r_{it}}{r_{nt}} \right) \ll z \mid \mathbf{p}, \theta \right], \text{ for } i = 2, \dots, n-1,$$

and define L_1 and L_n in a similar fashion.¹¹ $L_i[z \mid \mathbf{p}, \theta]$ represents the mass of agents with indifference ratios associated to security i lower than z , i.e., $r_i/r_j < z_j$ for $j \neq i$. For instance, $L_i[(1, \dots, 1) \mid \mathbf{p}, \theta]$ is the mass of agents that strictly prefer i over any other security when expected returns are equal across securities. In this context, we assume that there is enough heterogeneity so that demand for security i is bounded away from 0 and 1 when $\tau = 0$ and $ER_i = 1$ for all $i = 1, \dots, n$. That is, there is minimal “liquidity” at fair prices regardless of the market characteristics (\mathbf{p}, θ) .

Assumption 1. [Idiosyncratic belief heterogeneity] $L_i[(1, \dots, 1) \mid \mathbf{p}, \theta]$ is bounded away from zero for $i = 1, \dots, n$.

It is worth emphasizing that this assumption allows for the support of L_i to converge to a single point as $p_i \rightarrow 0$, as long as all the mass is not concentrated at $(1, \dots, 1)$ when $p_i > 0$.¹² In other words, it allows for subjective beliefs to be absolutely continuous with respect to true probabilities, and it does not implicitly require the distribution of subjective beliefs to exhibit “fat tails.” In addition, as we show in the Appendix, *Assumption 1* can be weakened when endowments and beliefs are independently distributed, by letting the distribution of beliefs exhibit “vanishing tails,” and our results would still go through.¹³

¹⁰If $ER_i > \max_{j \neq i} ER_j$ for some i then all traders invest in security j , leading to $\rho_i = 1$ and $\rho_j = 0$ for all $j \neq i$, implying $ER_j > ER_i$, a contradiction.

¹¹ $L_1[z \mid \mathbf{p}, \theta] = P \left[\left(\frac{r_{1t}}{r_{2t}}, \dots, \frac{r_{1t}}{r_{nt}} \right) \ll z \mid \mathbf{p}, \theta \right]$, and $L_n[z \mid \mathbf{p}, \theta] = P \left[\left(\frac{r_{nt}}{r_{1t}}, \dots, \frac{r_{nt}}{r_{n-1t}} \right) \ll z \mid \mathbf{p}, \theta \right]$.

¹²For instance, L_i could have full support in $[1 - \varepsilon, 1 + \varepsilon]^{n-1}$ with $\varepsilon \rightarrow 0$ as $p_i \rightarrow 0$.

¹³Independence of beliefs and endowments guarantees that heterogeneity in investment decisions does not vanish faster than heterogeneity in beliefs.

The next result shows that heterogeneity induces the FLB, even in the presence of transaction costs. In particular, any security i with a sufficiently high value (p_i) is underpriced in equilibrium, while the remaining (low value) securities are overpriced on average.

Theorem 1. [FLB] *If Assumption 1 holds, there exists $\bar{\tau} > 0$ such that for all $\tau < \bar{\tau}$ a necessary consequence of equilibrium is that there exists $\bar{q} < 1$ such that, for all $i = 1, \dots, n$, if $p_i > \bar{q}$ then security i is underpriced while securities $j \neq i$ are overpriced on average.*

Proof. In the Appendix □

Why does the FLB arise when agents are heterogeneous? Three key features in our model help explain it. On the institutional side, the underlying value of an Arrow-Debreu security is bounded, since it is given by its payout probability.¹⁴ On the agent side, we assume that agents have finite endowments and are risk neutral. Hence, they focus on relative comparisons of (subjective) expected returns across securities, choosing to invest their endowment on the highest return security. Finally, we assume that belief heterogeneity is idiosyncratic, i.e. represented by some dispersion of subjective expected returns around the true underlying expected returns. In this context, risk neutrality ensures that demand for each security does not vanish when true expected returns are equal across securities, regardless of payout probabilities. As a consequence, the presence of this minimal demand pushes the price of a security above its payout probability whenever the latter is very small, causing the security to be overpriced.¹⁵

We provide some intuition by focusing on the case of two securities, unit endowments ($w_t = 1$ for all $t \in T = [0, 1]$) and zero transaction costs. Note that, in this context, if the returns across the securities are (ER_1, ER_2) , then the share s_1 of investment in security 1 is bounded above and below by

$$L_1[ER_1/ER_2 \mid \mathbf{p}, \theta] \leq s_1 \leq 1 - L_2[ER_2/ER_1 \mid \mathbf{p}, \theta].$$

These bounds are tight because agents with $r_{1t}/r_{2t} = ER_1/ER_2$ are indifferent towards any investment (x_1, x_2) with $0 \leq x_1 + x_2 \leq 1$.

In this market idiosyncratic heterogeneity implies the existence of a positive lower bound of demand for security 1 when $ER_1 = ER_2$, given by $\underline{q} = \inf\{L_1[1 \mid \mathbf{p}, \theta] : (\mathbf{p}, \theta) \in \text{int}\Delta \times \Theta\}$. But then, when p_1 is sufficiently low, this lower bound \underline{q} on demand prevents the share of security i and thus its price to fall below p_1 , leading to overpricing of security

¹⁴This is a common feature of many asset markets. Many derivative contracts, such as the now notorious credit default swaps, are in essence contingent claims specifying a predetermined payout triggered by the occurrence of some event.

¹⁵The underpricing of securities with a high payout probability follows the same argument by noting that prices in the market add up to one.

1 ($ER_1 < 1$) and underpricing of security 2 ($ER_2 > 1$ since $p_2 = 1 - p_1 < 1 - s_1 = s_2$). More formally, let $p_1 < \underline{q}$ and suppose the theorem does not hold, i.e., $ER_1 \geq ER_2$. Then, by the above bounds we must have $s_1 \geq L_1[1 \mid \mathbf{p}, \theta] \geq \underline{q} > p_1$. But this is a contradiction since $s_1 > p_1$ implies that $ER_1 < ER_2$.¹⁶ A similar intuition applies when p_1 is sufficiently high, given the upper bound $1 - L_2$.

Remark. It is worthwhile to note the generality of our condition. Idiosyncratic heterogeneity only requires that there is some dispersion of beliefs around the correct beliefs, i.e. that there are minimal gains from trade in each security at fair prices. The result is thus a considerable generalization of existing models relating heterogeneity in beliefs to the longshot bias, such as Ali (1977).¹⁷

Example. *Suppose there are two assets in the market with probability of paying out a dollar being .9% and 99.1%, respectively, and there are 99 risk neutral agents with accurate beliefs and only one risk neutral agent who believes that the probability of payout on the longshot is anything better than 1%. If all agents have unit endowments, then the trading equilibrium involves this single agent investing in the longshot while all the remaining agents invest in the favorite.¹⁸ Thus a very small departure in beliefs of one percent of the population creates a sizable disparity in returns: the expected return on the longshot in the heterogeneous agent economy is -10 percent whereas the favorite has a positive expected return.*

4 Data

Betting markets, and racetrack betting in particular, are textbook illustrations of Arrow-Debreu securities markets. We focus our attention on the “win odds” market, which is the market for bets on which horse will win. This is considered the most competitive market at the racetrack, given it has the most liquid pool of money. Not surprisingly, it has been the subject of the most empirical attention.

Prices at the racetrack are quoted in terms of the odds R_i on horse or security $i = 1, \dots, n$, which are defined as the *net* return per dollar bet on security i in the event i wins the race—the *gross* return is given by $R_i + 1$. In North American racetracks, the odds are determined through a “parimutuel” system of wagering in which the losers pay

¹⁶Notice that all that we need is $L_1[1 \mid \mathbf{p}, \theta] > \alpha p_1$ for some $\alpha > 1$ and all p_i small enough. Hence, we can replace the uniform bound on L_1 with this “vanishing bound,” as long as $s_1 \geq L_1$ holds, which is the case when endowments and beliefs are independent. See Appendix A.1 for details.

¹⁷Ali’s model of a betting market considers two securities and assumes that $Pr(r_{1t}/r_{2t} \leq 1 \mid \mathbf{p}, \theta) = .5$ for any (\mathbf{p}, θ) , i.e. the canonical belief is the median belief in the population.

¹⁸In this case, the agent with optimistic beliefs invests in asset 1, and all other agents invest in asset 2, then the A-D price is .01 for asset 1 and .99 for asset 2. Hence, the optimistic agent views the longshot as underpriced and the remaining agents view the favorite as underpriced, i.e., there are gains from trade.

the winners. This system ensures that there are no differences in liquidity across securities and is equivalent to the market clearing condition in an Arrow-Debreu securities market where shares equal prices. Accordingly, market odds satisfy

$$R_i = \frac{1 - \tau}{s_i} - 1 \quad i = 1, \dots, n. \quad (2)$$

Given a vector of observed odds (R_1, \dots, R_n) , we can invert (2) to recover the underlying A-D prices $\rho_i = s_i$ for $i = 1, \dots, n$.

The FLB is in fact a widely documented empirical pattern of returns across bets at the racetrack. To illustrate it, consider a large data set consisting of horse starts, i.e., a sample of horses that competed in some race. Each observation i corresponds to a pair $(win_i, R_i + 1)$, where $R_i + 1$ is the gross return that A-D security i pays conditional on horse i winning and win_i is an indicator variable for whether horse i won the race or not. Define A_i as the ex-post gross return on security i , i.e.,

$$A_i = \begin{cases} R_i + 1 & \text{if } win_i = 1 \\ 0 & \text{if } win_i = 0, \end{cases}$$

and thus the regression

$$E[A_i | \log s_i] = E[(R_i + 1) p_i | \log s_i] = f(\log s_i) \quad (3)$$


measures the expected gross return among securities that have the same log price ($\log s_i$). Using a sample of 176,652 races that were collected from North American tracks over 2003-2006,¹⁹ consisting of 1,456,512 horse starts, we estimate f by non-parametrically regressing A_i on $\log s_i$ using a locally linear kernel weighted regression.²⁰ *Figure 1* shows the estimated returns —along with their 95% confidence interval, which clearly exhibit the FLB: low value securities have lower expected returns, where the effect is particularly pronounced when comparing the extremes. The returns in *Figure 1* are net of the track take. Thus, since the average track take is about 0.19, these returns imply that, roughly, horses with prices below 0.2 are overpriced and those with prices higher than 0.2 are underpriced.

To understand the argument implicitly put forth by Thaler and Ziemba (1988) as to why the FLB as found in the data is so puzzling, let us recall Lemma 1 and consider the indifference ratios of a representative agent that clears the market in equilibrium. In order to clear the market, the agent must be indifferent across securities to sustain

¹⁹The data were collected through the efforts of members at the website paceadvantage.com. This data set is one of the largest to have been assembled to study betting behavior.

²⁰We use the Epanechnikov kernel and a rule of thumb bandwidth estimator.

Figure 1: Favorite Longshot Bias




positive asset demand (a necessary condition for equilibrium). The expected gross return to an average bet in the data is $ER_{avg} = .76$. We can thus use *Figure 1* to recover the representative agent's indifference ratio ($\frac{\bar{r}}{r\rho}$), which is the ratio associated to the 'average' security (representing a random bet) and a security with price ρ . This ratio represents the amount that the average security must pay in expected returns relative to the a security with price ρ in order for the agent to be indifferent, and we can back out this relationship from the data because the representative agent is indifferent across all securities for the market to clear. As *Figure 2* shows, the representative agent overvalues securities with prices .01 or less by over 50 percent, and securities with prices .005 or less by approximately 100 percent, compared to the average bet.

5 Estimating Heterogeneity

We now show how to estimate a model of an A-D economy where agents have heterogeneous beliefs. The goal is to estimate the pattern of belief heterogeneity that is consistent with both the betting data and the equilibrium constraints of the model. That is, we seek to understand whether a heterogeneous beliefs model can explain the observed FLB and, furthermore, what the estimated pattern of heterogeneity looks like. We will contrast this typical agent with the representative agent that rationalizes the data.

In order to illustrate our empirical strategy we first describe the general maximum

Figure 2: Representative Agent Indifference Ratios



likelihood (ML) approach for estimating preferences from aggregate betting data, which was first introduced by Jullien and Salanié (2000) in the context of a representative agent framework. We then show how to extend their framework to a population of heterogeneous agents. Let us recall that the dataset consists of a sample of K markets or races, which are assumed to be independent of each other. Each market $k = 1, \dots, K$ is defined by the number of securities n^k , a vector of odds $\mathbf{R}^k = (R_1^k, \dots, R_{n^k}^k)$, and the identity of the winner i_{win}^k . The data thus identify the empirical relationship between prices and fundamentals, i.e., for any number of securities n with odds (R_1, \dots, R_n) we can identify the underlying probabilities of winning $\mathbf{p}(R_1, \dots, R_n)$: intuitively, $p_i(R_1, \dots, R_n)$ is identified by the fraction of times horse i wins in the subset of races characterized by (R_1, \dots, R_n) .²¹

A key insight in Jullien and Salanié (2000) (which we shall abbreviate as JS) is to recognize that an equilibrium model of the betting market, described by parameters $\theta \in \Theta$, implies a relationship between prices and fundamentals, denoted $\mathbf{p}(R_1, \dots, R_n; \theta)$. The empirical strategy implicitly underlying JS is to find the parameters $\theta \in \Theta$ whose prediction about the equilibrium relationship between prices and fundamentals best matches the actual empirical relationship. This strategy is naturally implemented using maximum likelihood, and can be described by two key steps:

1. Given a choice of model parameters $\theta \in \Theta$, for each market k , find the vector of

²¹Formally, this can be identified by the non-parametric regression of the indicator variable that a horse i wins on the vector of odds in the race

unobserved payout probabilities $\mathbf{p}^k(\theta) = (p_1^k(\theta), \dots, p_n^k(\theta))$ that is consistent with *observed* odds \mathbf{R}^k under equilibrium. That is, we solve

$$\mathbf{p}^k(\theta) = \phi(\mathbf{R}^k, \theta), \quad k = 1, \dots, K, \quad (4)$$

where ϕ is a model-specific mapping from odds to probabilities.

2. Estimate θ by maximizing the log-likelihood

$$LL(\theta) = \sum_{k=1}^K \log p_{i_{win}^k}^k(\theta). \quad (5)$$

This ML estimator consistently estimates the value of $\theta_0 \in \Theta$ that minimizes the Kullback-Leibler distance between the model $\mathbf{p}(R_1, \dots, R_n; \theta_0)$ and the data $\mathbf{p}(R_1, \dots, R_n)$.²²

JS formally showed that, in a representative agent framework of market equilibrium, there exists a unique equilibrium mapping between prices and fundamentals that is required for Step 1. In particular, the equilibrium mapping is determined by the set of indifference conditions

$$U(p_1^k, R_1^k; \theta) = U(p_2^k, R_2^k; \theta) = \dots = U(p_J^k, R_J^k; \theta), \quad (6)$$

where $U(p_i^k, R_i^k; \theta)$ is the payoff to the agent from investing in a security with payout probability p_i^k at odds R_i^k . That is, \mathbf{p}^k is solved to make the agent indifferent about investing across securities at the observed odds.²³ Using these conditions, preference parameters θ , such as risk aversion or cumulative prospect theory (CPT) coefficients, can be estimated. Using this approach, JS compared different representative agent preferences, and found evidence in favor of CPT over risk loving as the preferred explanation of the FLB.

A heterogeneous agents model, however, cannot be approached using this representative agent framework: the set of agents who are indifferent across all securities in a heterogeneous population typically has measure zero. Thus we can no longer use the equilibrium conditions (6) to solve for the relationship between prices and probabilities $\mathbf{p}^k(\theta) = \phi(\mathbf{R}^k, \theta)$ in a heterogeneous agents model (where the parameter $\theta \in \Theta$ determines the pattern of heterogeneity in the underlying betting population). Instead, the additional information that we will use to solve for the equilibrium relationship between prices and fundamentals are the market shares of the securities $\mathbf{s}^k = (s_1^k, \dots, s_n^k)$, which are derived from observed odds using (2). Note that market shares are not needed in the representative

²²Standard errors follow from standard likelihood theory, which JS give explicitly for this context in Section 4.c of their paper.

²³Jullien and Salanié (2000) show that, if U is continuous and respects first order stochastic dominance, the mapping ϕ from odds to probabilities implied by (6) is well-defined.

agent approach. However a heterogeneous agents model predicts market shares in a natural way. In particular, if U_t denotes the payoff function of agent t and we assume that she invests her endowment on a single security,²⁴ market shares are given by the aggregation of individual investments:

$$s_i^k = \frac{1}{\int_T w_t dt} \int_T w_t \mathbf{1} \left[U_t(p_i^k, R_i^k) > U_t(p_j^k, R_j^k) \forall j \neq i \right] dt, \quad i = 1, \dots, n^k, \quad (7)$$

where $\mathbf{1}[\cdot]$ is the indicator function. Thus, our goal is to characterize U_t in terms of our model of belief heterogeneity and show that the system of equations (7) can be uniquely inverted to recover \mathbf{p}^k . In order to be consistent with our theoretical model and to ease exposition, we focus here on the case of a population of risk neutral agents. Nonetheless, we describe in the Appendix the general identification strategy when the trader population exhibits heterogeneity in both risk preferences and beliefs.²⁵

To characterize (7) under risk neutrality and heterogeneous beliefs, let us return to *Lemma 1*, which allows us to represent agent preferences in terms of their relative deviations from correct beliefs $\mathbf{r}_t = (r_{1t}, \dots, r_{nt}) > 0$ for $t \in T$. Recall that agent t 's *subjective* expected returns of security i can be written as $\frac{ER_i}{r_{it}}$. Because choices are invariant to a monotonic transformation of preferences, we can express her utility as

$$U_t(p_i^k, R_i^k) = \log ER_i^k + \nu_{it}. \quad (8)$$

where we let $\nu_t = (\nu_{1t}, \dots, \nu_{nt})$ and $\nu_{it} = -\log r_{it}$. Abusing slightly our earlier notation, let $P(\cdot; \theta)$ be the distribution of ν_t , which captures the distribution of belief heterogeneity in the population T and is indexed by an unknown vector of parameters $\theta \in \Theta$ that is the object of estimation. Assuming that P is continuous and that the distribution of endowments is independent of P , by *Lemma 1* we can express (7) as

$$s_i^k = \int_{\nu_t \in \mathbb{R}^n} \mathbf{1} [\delta_i + \nu_{it} > \delta_j + \nu_{jt} \forall j \neq i] dP(\nu_{1t}, \dots, \nu_{n^k t}; \theta) \quad i = 1, \dots, n^k, \quad (9)$$

where $\delta_i = \log ER_i^k - \log ER_1^k$.

It turns out that this system of equations (9) is isomorphic to a differentiated product market in which each agent chooses among n differentiated products, with δ_i representing

²⁴We are implicitly modelling the decision of which asset to invest and abstracting from the participation decision since we do not have data on the number consumers who do not participate in the market. Rather, our data only gives us demand conditional on participation. In addition, we assume that the probability of a positive mass of agents being indifferent between two securities is zero. This is the case when the distribution of beliefs is continuous, as we assume below.

²⁵When we add risk aversion to the heterogeneous beliefs model, we find it does not empirically add much and does not affect our main findings as will be later discussed.

the “mean utility” of product i —the mean utility of product 1 is normalized to be zero— and $\nu_t = (\nu_{1t}, \dots, \nu_{n^k t})$ is a vector of random utility terms that is heterogeneous in the population. Inverting systems of equations defined by (9) is central to the demand estimation framework put forth by Berry et al. (1995) (aka BLP), i.e., solving for mean utilities given the underlying distribution of preferences over a set of differentiated products using the observed prices and market shares over these products. This connection provides us an immediately useful result: so long as P is continuous, then Berry (1994) shows that there exists a vector of mean utilities $(\delta_1, \dots, \delta_n)$ that solves (9). Furthermore, the more recent Berry et al. (2011) shows that such solution is unique.

We can now exploit these results to carry out the empirical strategy of JS in a heterogeneous agents context. Constructing the likelihood for given choice of parameters θ can be described by the following steps.

- (a) For any market k , we numerically find the unique underlying vector of mean utilities $\delta^k = (\delta_1^k, \dots, \delta_{n^k}^k) \in \mathbb{R}^n$ that solves the system of equations (9).
- (b) Given $\delta^k = (\delta_1^k, \dots, \delta_{n^k}^k)$, we can recover the underlying probability distribution over states of nature \mathbf{p}^k by using the following facts
 - (i) $\delta_i^k = \log ER_i^k - \log ER_1^k$,
 - (ii) The expected gross return is $ER_i^k = (1 - \tau) \frac{p_i^k}{s_i^k}$,
 - (iii) Probabilities $(p_1^k, \dots, p_{n^k}^k)$ over the n^k possible states of nature sum up to one.

6 Results

The only aspect of the estimation that remains to be discussed is how to parameterize the distribution $P(\nu_t; \theta)$. We wish to introduce belief heterogeneity so that it is *idiosyncratic*. To do so, we restrict P to satisfy $E[\nu_{it}] = 0$ for all $i = 1, \dots, n$, and $E(\nu_{it} | \nu_{-i,t}) = 0$ for any $\nu_{-i,t} \in \mathbb{R}^{n-1}$. That is, there is nothing systematic about the shocks ν_t . If we further assume that ν_{it} are i.i.d. $P(\nu_t; \theta)$ is non-parametrically identified, as we show in the Appendix.

The standard approach in the estimation of discrete choice demand models is to assume that ν_{it} is an i.i.d. “logit” error with common variance across the securities $i = 1, \dots, n$. That is, $\nu_{it} \sim F(\cdot | \sigma)$ where $F(\cdot | \sigma)$ is the CDF of an extreme value random variable with scale parameter σ . We now introduce a more flexible approach to modelling heterogeneity that applies to discrete choice models in general, but is particularly useful as a means of modelling belief heterogeneity. Instead of requiring that ν_{it} be a logit error, we let ν_{it} be

a *variance mixture* of logit errors. That is

$$P(\nu_{1t}, \dots, \nu_{nt}) = \int \prod_{i=1}^n F(\nu_{it} | \sigma) dG(\sigma),$$

where $G(\sigma)$ is the mixing distribution. Such a variance mixture retains the above key properties of idiosyncratic heterogeneity.²⁶ Moreover, as we argue in the Appendix (see section D), we can identify G given that P is (non-parametrically) identified. Solving for the inner integral in (9) analytically (see e.g., Train (2003)) yields the well known mixed logit demand

$$s_i = \int_{\sigma \in \mathbb{R}_+} \frac{\exp\left(\frac{1}{\sigma} \delta_i\right)}{\sum_{j=1}^n \exp\left(\frac{1}{\sigma} \delta_j\right)} dG(\sigma) \quad i = 1, \dots, n, \quad (10)$$

where the mixing takes place over the distribution of the variances σ . Although this system (10) bears a resemblance to the mixed logit demand that was originally estimated by Berry et al. (1995), there is a critical difference. Because our heterogeneity is governed by a variance mixture of the logit errors (as opposed to the usual mean mixture), it introduces a random coefficient $\frac{1}{\sigma}$ on the mean utility term δ_i . This random coefficient does not affect the existence and uniqueness of a solution to the system, but it does affect our ability to compute it because the contraction mapping proposed in Berry et al. (1995) is no longer valid (the mean utilities cannot interact with random coefficients in their setup). Instead, we minimize the sum of squared errors between observed shares and predicted shares using a quasi-newton procedure.²⁷

A one component mixture, i.e., a distribution $G(\sigma)$ with only one point in its support, corresponds to the standard logit assumption. Adding components to the simple logit gives rise to a finite mixture. We view this as a natural way to capture heterogeneity in the population in our context as the different components can be interpreted as trader “types”. A finite mixture with J components corresponds to a parameter vector $\theta = (\sigma_1, \dots, \sigma_J; P_1, \dots, P_J)$, where σ_j for $j = 1, \dots, J$ are the support points of the finite mixture and P_j is the probability mass of component j .²⁸ The parameter estimates for a one, two and three component mixture are presented in *Table 1*.

The first thing to note is that the two-component specification provides a much better

²⁶Conditional on σ , $F(\nu_{it} | \sigma)$ represents i.i.d. logit errors with common variance σ and thus satisfies idiosyncratic heterogeneity. Then, by integrating out over σ , idiosyncratic heterogeneity is retained.

²⁷The reason this computational strategy works well is that we have sensible starting values, namely the vector of zero mean utilities, which corresponds to all horses in the race having same expected return. Since these starting values are close to the true inverse when payout probabilities are not too extreme, the quasi-newton procedure is able to deliver a fast solution.

²⁸Any distribution over a compact support can be approximated by a multinomial distribution. See e.g., *Theorem 10.15* in Aliprantis and Border (2006).

Table 1: Parameter Estimates

Parameter	One Type Estimate (std. error)	Two Types Estimate (std. error)	Three Types Estimate (std. error)	Two Types-CARA Estimate (std. error)
σ_1	0.067 (0.0035)	0.028 (0.0033)	0.014 (0.0058)	0.034 (0.0074)
σ_2	-	0.503 (0.0588)	0.690 (0.1757)	0.514 (0.0862)
σ_3	-	-	0.075 (0.0306)	-
P_1	1	0.716	0.461 (0.1237)	0.706
P_2	-	0.284 (0.0184)	0.219 (0.0390)	0.294 (0.020)
P_3	-	-	0.320	-
CARA(γ)	-	-	-	0.003 (0.0031)
Log-likelihood	-307,391.5	-307,291.8	-307,291.3	-307,288.0
LR Test ^a (p-value)	199.4 (<0.0001)	-	1.0 (0.6065)	7.6 (0.0058)


^aThe LR test statistic is given by $2 \ln LL_a - 2 \ln LL_0$, where LL_0 and LL_a represent the log-likelihood of the null and the alternative model, respectively (where the alternative model nests the null).

fit that the one-component specification. To see this fact, we illustrate in *Figure 3* the pattern of (net) expected returns implied by the one-component and the two-component specifications, respectively.²⁹ While the one component model does a poor job fitting the expected returns of longshots, the two component model fits the observed pattern remarkably well. This improvement is further confirmed by the likelihood ratio test, which strongly rejects the one-component in favor of the two-component model.

Table 1 also shows that adding a third component does not improve the fit: the log likelihood is almost identical to the two-component model—the LR test does not find a significant difference in the likelihood associated with the two-component and three-component models. As can be seen, the introduction of a third component roughly reproduces the two

²⁹The x-axis in each figure is drawn on a log-scale. We can produce the model’s analogue of *Figure 1* by regressing the predicted ER_i^k on $\log s_i^k$ for all securities i and markets k in the data.

Figure 3: Predicted and Actual Returns



component case, in which the predominant component ($P_1 = 0.716$) exhibits low variance while the second component exhibits substantial dispersion. Consequently, in what follows, we focus on the two-type model.

In order to check our assumption of risk neutrality, we also estimate a two-component model in which agents have (homogeneous) CARA preferences following the general approach laid out in the Appendix. The last column of *Table 1* shows the model estimates. The estimated CARA coefficient (denoted by γ) yields a tiny degree of risk aversion ($\gamma = 0.003$) albeit not significantly different from zero. That is, we fail to reject risk neutrality ($\gamma = 0$). In addition, the remaining parameters are virtually identical to the two-component risk neutral version of the model.³⁰

6.1 Economic Interpretation: Informed and Noise Traders

The two component mixture suggests a bimodality in the underlying distribution of agents in the market. In the context of financial markets, there is a large literature that theoretically introduces a bimodal distribution of traders. On the one hand there are the *arbitrageurs* who are motivated by revenue maximization and are relatively well informed. On the other hand there are the noise or liquidity traders (known as recreational gamblers

³⁰Nonetheless, allowing for CARA preferences yields a significant likelihood gain—the LR test rejects the risk neutral version of the model in favor of the CARA specification at the 1% level. However the economic implications of these two models are virtually identical.

in betting markets) who typically have a wider dispersion of beliefs, due to having less precise information than the arbitrageurs or to not properly updating their beliefs after observing prices. This two type classification is commonplace in the market microstructure literature (see e.g. Glosten and Milgrom (1985), Kyle (1985) and the ensuing literature), which separates traders into “noise” traders and insider speculators, in the ‘limits of arbitrage’ literature (De Long et al., 1990; Stein, 2009), and in research related to mispricing in speculative markets (Shin, 1991, 1992; Serrano-Padial, 2012). In the context of parimutuel markets, Ottaviani and Sorensen (2006) interpret these types as *late* (more informed and neutral in tastes) and *early* traders (less informed and more idiosyncratic in tastes).³¹

Our estimates are consistent with this bimodal view of the trading population. If we interpret the components of the mixture as representing different “types” of economic agents, then type σ_1 represents informed traders, which is also the modal type in the population (72%), while type σ_2 is a noise trader, since it exhibits a higher level of “noise” in subjective valuations. This is, to our knowledge, the first paper empirically identifying the existence of informed and noise traders with aggregate data on prices and asset value realizations.³²

To quantify the degree of ‘noise’ exhibited by type- σ_2 agents, consider the choice between investing in a security at random, which yields average returns in expectation, and an extreme longshot. For example, in our dataset, the average (gross) return among all assets is 0.76, whereas it is about 0.35 for securities in the bottom percentile of the price distribution. The probability that a type- σ_2 agent in our estimated population prefers a security with average returns to a security i yielding 0.35 is given by

$$Pr(0.76/0.35 > \bar{r}_t/r_{it}) = Pr\left(\bar{\nu}_t - \nu_{it} < \frac{(\log 0.76 - \log 0.35)}{0.503}\right) \approx 0.82,$$

where the last equality following from the fact that $\bar{\nu}_t - \nu_{it}$ is distributed standard logistic.³³

Furthermore, since type σ_2 only has a mass of approximately 0.28 in the population, and type σ_1 prefer the favorite approximately virtually 100 percent of the time, the probability of the random investment being preferred to the longshot in the population as a whole is approximately $0.82 \times 0.28 + 0.72 \approx 0.95$. This is in stark contrast with a representative agent, who must indifferent between the two bets in order to explain the data.

We further compare the representative agent and traders in our estimated model by looking at the distribution of indifference ratios $\frac{\bar{r}}{r_\rho}$ in the two-component population. *Fig-*

³¹Consistent with this two-type hypothesis, Forsythe et al. (1992) and Cowgill et al. (2009) find empirical evidence of two types using agent-level data from prediction markets.

³²There are several empirical studies that estimate the presence of noise traders using individual trading data from laboratory markets (Forsythe et al., 1992; Cipriani and Guarino, 2005), or transaction data from financial markets (Easley et al., 1997).

³³The difference of standard Type 1 extreme value random variables is standard logistic.

Figure 4 shows the different quantiles of the distribution, which we compare alongside the implied indifference ratios of a representative agent. Agents in the interquartile (IRQ) range, i.e., the inner 50% of the agents in the population, exhibit very little dispersion from the canonical agent, i.e., the agent with indifference ratios equal to 1.³⁴ That is, the modal behavior in the heterogeneous population is closely captured by a risk neutral agent with correct beliefs. This is driven by the fact that the first component has a very low variance. Beliefs become more dispersed for agents outside the IRQ, reflecting the fact that the variance of the second component is much higher ($\sigma_2 = 0.503$ versus $\sigma_1 = 0.028$). It is only at the extreme tails of our estimated distribution of belief heterogeneity that we approach the extreme beliefs of the representative agent with respect to longshots.

Figure 4: Individual vs. Market Behavior

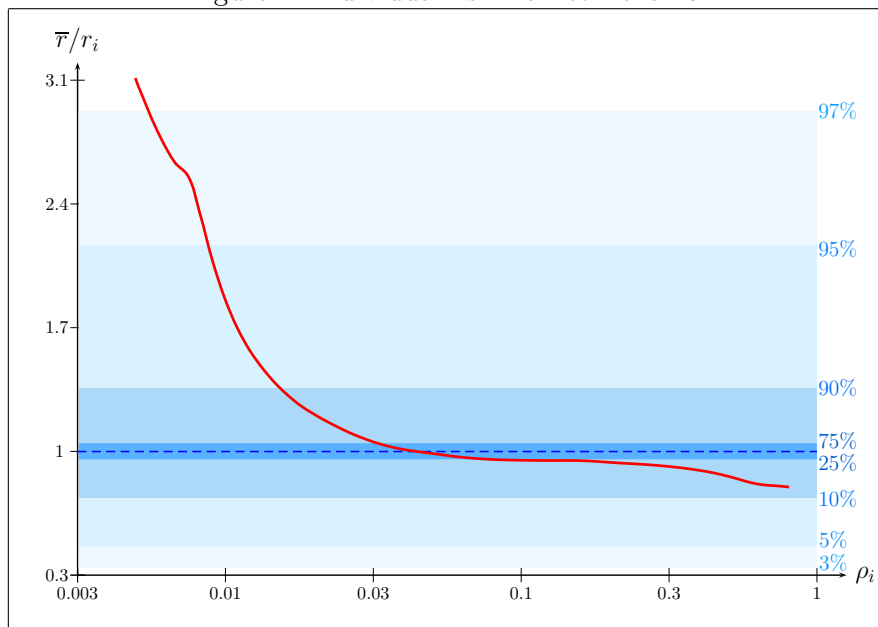


Figure 4 also provides intuition for why the two-type specification performs better than the one-type. Consistent with the numerical example in Section 3, to generate the observed disparity in returns between longshots and the average bet, we need only a small fraction of traders investing in the longshot. This is because the empirical winning probabilities of longshots are very small. Accordingly, most of the agents belonging to the first component exhibit beliefs that induce them to invest on the favorites, while agents in the second component spread their investment across securities, with a minority of them going for the extreme longshots. In contrast, the one-type specification faces a trade-off: either it

³⁴The IRQ is a standard measure of dispersion and model behavior. Manski (2004) also uses it to document dispersion in beliefs.

exhibits low belief dispersion (as it is the case) thus fitting well the returns on securities with moderate and high payout probabilities but does not generate enough demand for the longshots, or it exhibits higher dispersion, hence fitting poorly the returns on favorites.

7 Preferences vs. Belief Heterogeneity

In this section we compare our belief heterogeneity model to the two leading explanations of the FLB: risk loving and probability overweighting. In particular, we estimate a model of heterogeneous risk preferences that generalizes the representative agent model of expected utility employed in JS. We also estimate a representative agent model of CPT preferences as was employed by JS. The three main empirical questions of interest are the following:

- Do most agents exhibit risk love in the heterogeneous risk preferences model (denoted HR for short)?
- Does the explanation of the FLB that emerged from JS’s estimation of CPT, namely the overweighting of loss probabilities, also arise in our data?³⁵
- How do each of these alternatives empirically fare against our belief heterogeneity model (henceforth BH)? Specifically, can we formally compare these three models and study whether one of them stands out as the better explanation of the data?

7.1 Heterogeneous Risk Preferences

We start by considering a model with heterogeneous risk preferences (as opposed to heterogeneity in beliefs) within an expected utility framework. As in JS, we assume that agents are expected utility maximizers with correct beliefs, common endowment w and initial wealth M , and make a discrete choice from the set of available securities. However, we depart from JS by allowing heterogeneous risk attitudes. Agent t ’s expected utility of investing in security i is given by $U(p_i, R_i, \gamma_t) = p_i u(M + wR_i, \gamma_t) + (1 - p_i)u(M - w, \gamma_t)$, where u represents utility over final wealth. The term γ is a utility parameter that is distributed according to a continuous distribution $F(\cdot; \theta)$ strictly increasing over the support of γ , where larger γ indicates more risk aversion. The parameter to be estimated is θ , which determines the pattern of heterogeneity in the population of agents.

Observe that if we order horses such that $p_1 < \dots < p_n$, then we must have $R_1 > \dots > R_n$ in any equilibrium with trade. This is driven by the fact that expected utility respects

³⁵This is of interest because their approach has not been applied beyond the UK betting markets they originally considered, which follow a fixed-odds wagering system rather than a parimutuel mechanism. As Ottaviani and Sorensen (2005) show, in the presence of asymmetric information, these differences in market structure could have significant pricing implications.

FOSD, i.e., agent t 's expected utility from investing in i is increasing in both R_i and p_i . Thus, if $(R_j, p_j) > (R_i, p_i)$ for some i, j then all traders, regardless of γ_t , will prefer j over i and the market will not clear.

This monotonicity of odds w.r.t. probabilities allows to solve for \mathbf{p}^k using \mathbf{R}^k for any parameter value $\theta \in \Theta$ and thus conduct step 1 in the above ML estimation. This is because, given (R_j, p_j) and (R_i, p_i) with $p_j > p_i$, if an agent with risk coefficient γ_t prefers j over i then all agents that are more risk averse ($\gamma_{t'} > \gamma_t$) also prefer j to i . From this single crossing condition it immediately follows, as the next result states, that the mass of agents investing on i is given by the agents with $\gamma_t \in [\gamma_{i-1}, \gamma_i]$, where γ_i corresponds to the *marginal* agent indifferent between security i and $i + 1$.³⁶

Theorem 2. *If $p_1 < \dots < p_n$ and $R_1 > \dots > R_n$ then there exist $-\infty = \gamma_0 < \gamma_1 < \dots < \gamma_{n-1} < \gamma_n = \infty$ such that the mass of agents investing in security i is given by $F(\gamma_i; \theta) - F(\gamma_{i-1}; \theta)$, and*

$$U(R_1, p_1, \gamma_1) = U(R_2, p_2, \gamma_1), \dots, U(R_{n-1}, p_{n-1}, \gamma_{n-1}) = U(R_n, p_n, \gamma_{n-1}). \quad (11)$$

This result implies that for any race k , the system of market shares equations (7) is given by

$$s_i^k = F(\gamma_i^k; \theta) - F(\gamma_{i-1}^k; \theta), \quad i = 1, \dots, n^k, \quad (12)$$

where $\{\gamma_i^k\}_{i=1}^{n^k}$ satisfy (11). Unlike the BH model, (12) admits a closed form solution. We describe in the Appendix the steps to analytically recover the underlying probabilities.

Following JS and Snowberg and Wolfers (2010), we estimate a model in which agents have CARA preferences, given by $u(x, \gamma) = (1 - e^{-\gamma x}) / \gamma$. A key feature of CARA utility is that the initial wealth M drops out of the indifference conditions and the endowment w gets subsumed in the definition of γ , i.e., γ and w are not separately identified and thus w is implicitly normalized to 1. We parameterize F as a normal distribution with mean μ_γ and standard deviation σ_γ .³⁷ Our estimates are shown in the first column of *Table 2*. As can be seen, virtually all agents in the HR model exhibit risk loving attitudes (98.8% of the population), with very little dispersion: γ ranges between -0.02 and 0. The risk attitudes we find are close to the representative agent estimates of Snowberg and Wolfers (2010) ($\gamma = -0.017$) and exhibit slightly less risk loving than in JS ($\gamma = -0.055$).

The intuition for why heterogeneous risk preferences do not lead to any fundamentally

³⁶For more details see Chiappori et al. (2009), who provide an overview of the general approach to identification and estimation of contract choice models under a single crossing condition on heterogeneity. They also suggest different applications from parimutuel markets to insurance and the statistical value of life.

³⁷We have tried alternative flexible distributions and the differences in both the estimated distribution of CARA coefficients and the likelihoods were insignificant.

Table 2: Preference-Based Models

Parameter	HR	CPT	JS (CPT)
	Estimate (std. error)	Estimate (std. error)	Estimate (std. error)
CARA (μ_γ, γ)	-0.0086 (0.0017)	-0.032 (0.0006)	-0.072 (0.021)
Std dev (σ_γ)	0.0038 (0.0012)	-	-
Gains (α)	-	1.220 (0.0003)	1.162 (0.143)
Losses (β)	-	0.280 (0.0006)	0.318 (0.272)
Observations	176,466	176,466	4,037
Log-likelihood	-307,333.6	-307,301.5	-7,365.3
Vuong Test (p-value)	370.70 (<0.0001)	34.92 (<0.0001)	-

new insights beyond the representative agent risk loving explanation is the following: if returns are increasing in probabilities (or if favorites yield higher returns than the other securities), as it is the case in the data, risk averse agents will only invest in the favorite. To see why, observe that if $ER_i = ER_j$ with $p_i > p_j$ then security j is a mean preserving spread of i and a risk averse agent would always prefer i to j . Accordingly, a positive mass of risk lovers is needed to generate demand for the remaining securities and for the FLB to emerge in equilibrium. Indeed, all the marginal agents in (11) must be risk lovers. Thus, allowing for the possibility of risk averters in the population does little to change the tight link between risk love and the FLB from a pure risk preferences perspective.

7.2 Cumulative Prospect Theory

The CPT model we estimate is the preferred specification in JS: a three-parameter model consisting of a single CARA utility function with risk aversion coefficient γ to capture utility over both monetary gains and losses, and two separate probability weighting functions: one for gains and the other for losses.³⁸ As in JS, we assume that the representative agent has

³⁸Notice that the JS specification is richer than the one-parameter rank-dependent utility model of Snowberg and Wolfers (2010).

correct beliefs about p . The value of investing w on security i given odds R_i is given by

$$U(p_i, R_i) = G(p_i)u(wR_i, \gamma) + H(1 - p_i)u(-w, \gamma),$$

where $u(x, \gamma)$ is a CARA utility, $G(p) = p^\alpha$ is the weighting function for gains and $H(p) = p^\beta$ is the weighting function for losses.³⁹ We estimate the model by maximum likelihood using indifference conditions (6) —see JS for details.

Our parameter estimates along with the CPT estimates from JS are presented in the last two columns of *Table 2*. As can be seen, both estimates are quite close, but because of the much larger size of our sample our estimates have a much higher precision. We find that preferences exhibit slight risk loving ($\gamma < 0$), a slightly convex weighting of gains ($\alpha > 1$) and a highly concave weighting of losses ($\beta \ll 1$). Our estimates reinforce the major empirical finding from JS: under CPT, the key primitive driving the overpricing of longshots is the overweighting of loss probabilities (particularly for small probabilities).

7.3 Testing the Models

We have considered three alternative explanations of the FLB and the natural question becomes whether any one model stands out as a superior explanation of the data. We now aim to formally compare these various theories, both against each other and also against the true relationship between prices and fundamentals in the data. To motivate the econometric foundation for comparing these non-nested models, observe that the likelihood from our BH model exceeds that of the HR and CPT models. This suggests that our beliefs approach exceeds the explanatory power of these preference based alternatives. We can actually formalize this comparison in light of the identification strategy presented in Section 5. Let $\phi(\mathbf{R})$ denote the *true* mapping between prices and fundamentals in the data, i.e., $\mathbf{p} = \phi(\mathbf{R})$, and denote $\phi^m(\mathbf{R}; \theta)$ the mapping between prices and fundamentals predicted by model $m \in \{BH, HR, CPT\}$ for $\theta \in \Theta^m$. If model m is properly specified, then

$$\phi(\mathbf{R}) = \phi^m(\mathbf{R}; \theta^m) \tag{13}$$

for some ‘true value’ of parameters $\theta^M \in \Theta^M$. A natural question is whether (13) is actually testable, i.e., whether we can test that model m is consistent with the underlying data generating process. The likelihood framework naturally allows for such a specification test, which was first proposed by White (1982). This specification test is derived from the information matrix equality, which is a fundamental theorem in likelihood theory (see e.g.,

³⁹Observe that a single initial wealth level does not have to be assumed here because the CPT model considers only the utility of gains and losses relative to existing wealth.

Cameron and Trivedi, 2005). In our current context, the information equality states that, if (13) holds, the following must be satisfied:

$$E \left[\frac{\partial^2}{\partial \theta \partial \theta'} \log \phi_{i_{win}}^m(\mathbf{R}; \theta^m) \right] = -E \left[\frac{\partial}{\partial \theta} \log \phi_{i_{win}}^m(\mathbf{R}; \theta^m) \frac{\partial}{\partial \theta} \log \phi_{i_{win}}^m(\mathbf{R}; \theta^m)' \right], \quad (14)$$

where the expectation is taken with respect to the true data generating process over winners and odds (i_{win}, \mathbf{R}) in the data. White's test uses the sample analogues of each side of the equality to construct a quadratic form that should be sufficiently close to zero if (13) is true. We use the form of the test given by White (1987). Specifically, for each race k in the data let $l_k(\hat{\theta}^m) = \log \left(\phi_{i_{win}}^m(\mathbf{R}^k; \hat{\theta}^m) \right)$, i.e., the log likelihood value associated with race k at the estimated parameters $\hat{\theta}^m$. The test statistic is constructed as

$$T^m = \left\{ K^{-1/2} \sum_{k=1}^K \hat{q}_k \right\}' \left\{ \frac{1}{K} \sum_{k=1}^K \hat{q}_k \hat{q}_k' \right\}^{-1} \left\{ K^{-1/2} \sum_{k=1}^K \hat{q}_k \right\},$$

where

$$\hat{q}_k = \text{vech} \left\{ \frac{\partial^2 l_k}{\partial \theta \partial \theta'}(\hat{\theta}^m) + \frac{\partial l_k}{\partial \theta} \frac{\partial l_k}{\partial \theta'}(\hat{\theta}^m) \right\}.^{40}$$

If (13) holds, T^m is asymptotically distributed χ^2 with $J(J+1)/2$ degrees of freedom, where J is the length of the parameter vector $\theta^m \in \Theta^m$. Hence "large" values of the test statistic T^m indicate a significant difference between the model's predictions and the data.

The results of the test are rather stark. The test statistic fails to reject (13) for the BH model at standard significance levels ($T^{BH} = 8.54$ whereas the 10% critical value is 10.64). Hence we cannot reject the null hypothesis that the BH model equals the true data generating process. In contrast, the test statistic for the CPT model blows up considerably: $T^{CPT} = 1,139.08$, which clearly rejects the null that the CPT model describes the true DGP. Hence, although the log-likelihood values of the two models are somewhat close (BH having a higher likelihood than CPT), as far as the information equality test we see a big difference between these models in terms of their ability to accurately describe the true data generating process. Not surprisingly, the HR model fares even worse, given that its likelihood is substantially lower than that of CPT. We don't interpret this test as literally telling us that the two-type BH model is entirely properly specified. However, we do interpret it as saying that the model captures the subtle variation in the data well enough such that it cannot be distinguished from the true data generating process with nearly 200,000 races and only 3 parameters!

We reach similar conclusions if instead we implement the pairwise non-nested model

⁴⁰vech(A) denotes the half-vectorization of symmetric matrix A .

selection tests proposed by Vuong (1989) between the BH and the HR and CPT models, which are shown at the bottom of *Table 2*.⁴¹ The Vuong test is somewhat weaker than the White test since it only tests which model is closer to the true DGP as opposed to testing whether the proposed model is equal to the true DGP, however it has a somewhat more transparent interpretation and thus it is useful to nevertheless consider it. Observe that both the CPT and the BH models have the same number of parameters. As can be seen there, we reject that either the HR or CPT models are as close to the true DGP as the BH model, which confirms the insight from the White test.

7.4 Information Differences and the FLB

We have shown that our estimated two-type model of belief heterogeneity better explains the FLB as compared to the HR and CPT models. Indeed we cannot reject the hypothesis that the BH model we estimate is the true data generating process underlying the relationship between prices and fundamentals, even with close to 200,000 races in the data. We now ask whether we can gain further confidence that it is indeed belief heterogeneity what accounts for the superior empirical performance of the random utility model (10) we estimate. That is, although we derived our random utility on the basis of belief heterogeneity, and our estimates of the model have a natural interpretation in terms of informed and noise traders, in principle, the heterogeneity we estimate may have other interpretations.⁴²

To speak to this issue, we exploit a form of variation in the data that allows us to even further distinguish between a beliefs versus preferences interpretation of the FLB. Specifically, because beliefs respond to information while preferences do not, an ideal experiment would be to use differences across races in the amount of information about participating horses. Unfortunately, our data does not contain any horse specific information. However, we can instead make use of the fact that horse races on the same day at the same track come in two different forms: *maiden* and *non-maiden* races. Horses in a maiden race are those that have yet to win a single race, and thus by definition any new horses are entered

⁴¹The Vuong test is based on testing the null hypothesis that

$$H_0 : E_0 \left[\log \frac{\phi_{i_{win}}(\mathbf{R}; \theta^m)}{\phi_{i_{win}}(\mathbf{R}; \theta^{m'})} \right] = 0,$$

where the expectation is taken with respect to the true data generating process over (i^{win}, \mathbf{R}) and θ^m and $\theta^{m'}$ denotes the pseudo-true value of the parameter vector under models m and m' respectively (which is what our MLE estimator can be shown to consistently estimate when the models are potentially misspecified as descriptions of the true DGP). This null hypothesis tests whether both models are equidistant (in the Kullback-Leibler sense) to the true DGP. Because our model's are completely non-nested, we can use the simple one-step implementation of the test described by Vuong that only requires computing the likelihood value corresponding to each observation. We also avoid the problem of size distortion that arises under the Vuong test when models are not completely non-nested, recently raised by Shi (2012).

⁴²For example, idiosyncratic tastes (unrelated to risk preferences) that arise from state-dependent utility.

into maiden races whereas horses must have a racing history to participate in a non-maiden race. Thus on average, there is much less handicapping information about horses in maiden races as compared to non-maiden (Camerer, 1998; Mitchell, 1989).

The greater uncertainty surrounding maiden races has testable implications for the theories we have explored. For a given vector of winning probabilities \mathbf{p} , the equilibrium predictions of a purely preference-based model would be identical for the two types of races because information has no bearing on demand. In contrast, one should expect less belief dispersion in non-maiden races since traders can condition their beliefs on richer (public) information, which in turn affects equilibrium prices—a fact we discuss in more detail below. Thus, unlike preference-based theories, a belief heterogeneity perspective predicts a change in the magnitude of the FLB across maiden and non-maiden races.

To compare these predictions, we control for other differences that might exist between maiden and non-maiden races by restricting our samples to *claiming* races, in which horses can be purchased before the races and exhibit horses with similar price tags (thus trying to ensure a level playing field). This is the most frequent type of races in the US (over 54% of all races). Furthermore, maiden and non-maiden claiming races take place typically on the same day at the same track,⁴³ and they exhibit similar track takes.⁴⁴ Accordingly, we should expect little differences in the trader population across the two types of races.

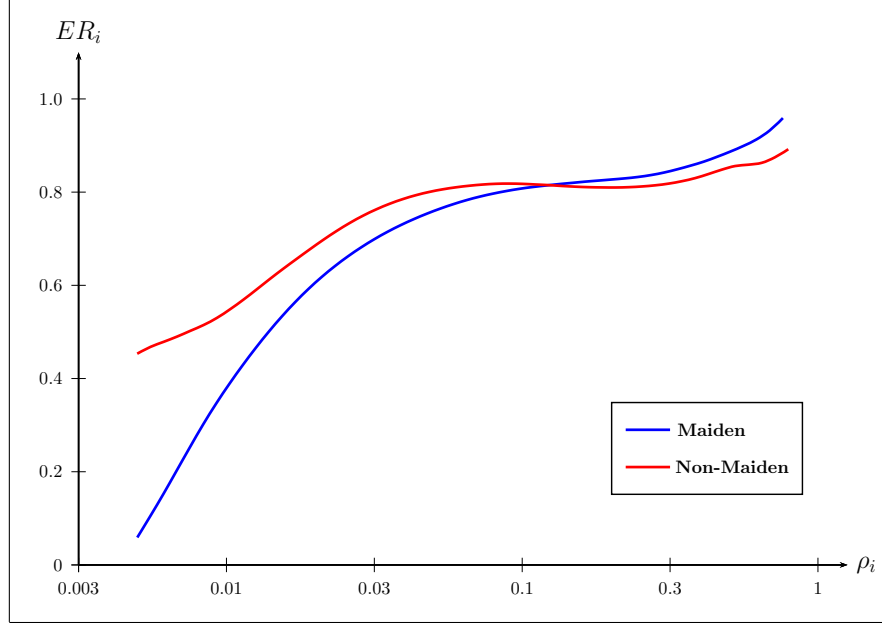
Figure 5 shows the relationship between prices and returns across race types. There is a considerably smaller disparity in returns across assets in non-maiden races as compared to maiden. That is, the FLB is much more pronounced in maiden races. Specifically, while expected returns in maiden races go from \$0.05 for extreme longshots to \$0.96 for heavy favorites, returns in maiden races are much more compressed, ranging from \$0.45 to \$0.89.

The substantial differences in the pattern of returns hint at a basic inconsistency of a preference-based view of the FLB: it can only explain pricing differences across types of races through differences in preferences. We illustrate this inconsistency by estimating the HR and CPT models on the separate subsamples. *Table 3* presents the HR estimates, which exhibit significant differences in the distribution of risk attitudes (we reject the mean γ is equal cross samples at the 1% level). Similarly to the full sample estimates, the average agent exhibits risk loving in maiden races while being essentially risk neutral in non-maiden races—the mean γ is actually positive (risk averse) but not significantly different from zero. *Figure 6* shows that the distribution of risk attitudes is slightly shifted toward risk aversion in non-maiden races while the vast majority of traders are risk lovers in maiden races. CPT estimates are given in *Table 4*, which show significant differences across

⁴³In our dataset, 99.5% of maiden races were run in a day where non-maiden races were also run.

⁴⁴The average track take for maiden claiming races is 0.190 with std.deviation 0.04. For non-maiden claiming races the average take is 0.192 and the sdt.dev. is 0.04.

Figure 5: The FLB in Maiden and non-Maiden Races



samples, in particular on the probability weighting of losses (the difference is significant at the 1% level): while both the full sample and maiden parameters are very similar, β is more than twice as high in non-maiden races, resulting in a much less concave weighting function (see *Figure 7*). These large swings in preferences across maiden and non-maiden races are difficult to rationalize.

Table 3: HR Estimates for non-Maiden and Maiden Races

Parameter	non-Maiden Estimate (std. error)	Maiden Estimate (std. error)	Full Sample Estimate (std. error)
CARA (γ)	0.0032 (0.0025)	-0.0089 (0.0039)	-0.0086 (0.0017)
Std dev (σ_γ)	0.0080 (0.0017)	0.0057 (0.0027)	0.0038 (0.0012)
Observations	87,394	29,003	176,466
Log-likelihood	-156,098.7	-50,518.8	-307,333.6

The more pronounced FLB in maiden races, on the other hand, is fairly natural under a belief heterogeneity view: with less information on the horses, we should expect more belief dispersion in maiden races, particularly among noise traders, which translates into a more pronounced FLB. However our beliefs model does still impose a restriction we can

Figure 6: Risk Heterogeneity in Maiden and non-Maiden Races





Table 4: CPT Estimates for non-Maiden and Maiden Races

	non-Maiden	Maiden	Full Sample
Parameter	Estimate (std. error)	Estimate (std. error)	Estimate (std. error)
CARA (γ)	-0.020 (0.0029)	-0.032 (0.0041)	-0.032 (0.0006)
Gains (α)	1.12 (0.0307)	1.20 (0.0420)	1.22 (0.0307)
Losses (β)	0.55 (0.0898)	0.25 (0.0769)	0.28 (0.0006)
Observations	87,394	29,003	176,466
Log-likelihood	-156,101.0	-50,508.8	-307,301.5

test: the change in the FLB across the two race types should be driven by a change in the belief dispersion of each type and *not* by a change in the proportion of trader types. If our model instead were to empirically require a higher prevalence of noise traders in the population to explain the change in the pattern of returns, it would suffer from the same parameter instability problem that the above preference theories suffer from. We formulate this claim as two testable hypotheses.

Hypothesis 1. *The fraction of types in the BH model is the same across maiden and non-maiden races.*

Figure 7: CPT Weighting Function for Losses



Hypothesis 2. *Belief dispersion in the BH model is higher in maiden races compared to non-maiden races.*

The basic insight behind *Hypothesis 2* is that, if agents are Bayesian, their beliefs should get closer as the amount of (public) information about the underlying state of the world accumulates over time. For instance, in the context of horse racing, if agents observe the performance of horses over time, they would eventually agree on the (true) winning probabilities. This notion of belief convergence is formally treated in the literature on merging of opinions (Blackwell and Dubins, 1962; Kalai and Lehrer, 1994; Lehrer and Smorodinsky, 1996; Gossner and Tomala, 2008).^{45,46} But, if beliefs get closer as more information becomes available, we should also expect a smaller FLB in non-maiden races. We theoretically show in the Appendix that this is the case in a model where the same A-D security market is repeated over time.⁴⁷

Table 5 presents the BH estimates for the two samples, which confirm *Hypothesis 1* and

⁴⁵A related literature deals with learning in competitive markets and convergence to rational expectations equilibria. Examples are Townsend (1978); Feldman (1987); Easley and Blume (1982) and Vives (1993).

⁴⁶Although the mentioned research on merging concerns the evolution of beliefs in the long run, Sandroni and Smorodinsky (1999) show that the speed of convergence can be quite fast. In addition, Gossner and Tomala (2008) provide bounds on average prediction errors.

⁴⁷Specifically, we show that, for given \mathbf{p} , the FLB gets mitigated overtime when agents can observe the history of ex-post returns and agree on the *direction* of the belief updating. The latter means that, if state i is realized in some period, all agents increase their subjective belief about p_i for the next period. The notion of belief consistency we employ is *ordinal*, in contrast to the stronger notion of concordant beliefs common in the literature (e.g. Ottaviani and Sorensen, 2010a; Milgrom and Stokey, 1982).

Table 5: BH Estimates for Claim and non-Claim Races

Parameter	non-Maiden Estimate (std. error)	Maiden Estimate (std. error)	Full Sample Estimate (std. error)
σ_1	0.0001 (<0.0001)	0.0366 (0.0079)	0.028 (0.0033)
σ_2	0.340 (0.0180)	0.6599 (0.1403)	0.503 (0.0588)
P_1	0.700	0.727	0.716
P_2	0.300 (0.0104)	0.273 (0.0321)	0.284 (0.0184)
Observations	87,394	29,003	176,466
Log-likelihood	-156,090.0	-50,505.8	-307,291.8

2. The fraction of type one is very stable around 70-73% —we fail to reject *Hypothesis 1* at any standard significance level. In contrast, belief dispersion estimates are significantly smaller (at the 1% level) in non-maiden races: σ_1 is virtually zero and σ_2 is about half its value in maiden races. Thus, remarkably, our heterogeneous beliefs approach explains the change in prices across information structures in the theoretically predicted way.

8 Conclusion

We have shown in this paper that allowing for belief heterogeneity in asset markets can reconcile model predictions with observed aggregate patterns without compromising the validity of standard behavioral assumptions such as weak risk aversion and expected utility. We have used Arrow-Debreu security markets as a general setting in which to illustrate these points. We have also showed that a belief-based model of trade in financial markets outperforms existing preference based explanations of the FLB on several dimensions. We think heterogeneous beliefs have the potential to play an important role in empirical work in other institutional settings, such as insurance and credit markets, and we hope this paper encourages future research in this area.

Appendix

A Omitted Proofs

Proof. [Proof of Lemma 1] Consider the maximization problem of agent t with endowment $w_t > 0$ and beliefs $(\pi_{1t}, \dots, \pi_{nt}) \gg 0$. Given the state of the world (p_1, \dots, p_n) and market prices $(\rho_1, \dots, \rho_n) \gg 0$, the agent solves

$$\begin{aligned} \max_{(x_1, \dots, x_n) \in \mathbb{R}_+^n} \quad & \sum_{i=1}^n x_i \left(\frac{\pi_{it}}{\rho_i} - 1 \right) + w_t \\ \text{s.t.} \quad & \sum_{i=1}^n x_i \leq w_t. \end{aligned}$$

The ratio π_{it}/ρ_i represents the *subjective* (expected) returns of security i , i.e. the return given agent's beliefs. If there is a security i such that

$$\frac{\pi_{it}/\rho_i}{\pi_{jt}/\rho_j} > 1, \quad \forall j \neq i, \quad (15)$$

then the solution to the agent's problem implies investing all the endowment w_t in security i . This is because (i) security i yields the strictly highest subjective returns among all securities; and (ii) security i 's subjective returns are strictly greater than one. Property (ii) comes from the fact that, since $\sum_i \pi_i = 1$ and $\sum_i s_i = 1$, we must have $\max_h \frac{\pi_{ht}}{s_h} > 1$ whenever $\frac{\pi_{ht}}{s_h} \neq \frac{\pi_{kt}}{s_k}$ for some h, k . This also means that, if the latter is true, the agent will invest all her endowment in the securities yielding $\max_h \pi_{ht}/s_h$, being indifferent about how much to invest on each of them. Finally, if $\frac{\pi_{ht}}{s_h} = \frac{\pi_{kt}}{s_k}$ for all h, k , then subjective returns are all equal to one, making the agent indifferent between investing any amount in $[0, w_t]$. To finish this case, notice that, by letting $r_{it} = p_i/\pi_{it}$, expression (15) becomes

$$\frac{ER_i}{ER_j} > \frac{r_{it}}{r_{jt}}, \quad \forall j \neq i.$$

□

Proof. [Proof of Theorem 1] We consider first the case of a market with zero transaction costs $\tau = 0$. Let $q_i(ER_1, \dots, ER_n; \mathbf{p}, \theta)$ represent the aggregate investment on asset i and denote $\mathbf{r}_{it} = \left(\frac{r_{it}}{r_{1t}}, \dots, \frac{r_{it}}{r_{i-1t}}, \frac{r_{it}}{r_{i+1t}}, \dots, \frac{r_{it}}{r_{nt}} \right)$. Given *Lemma 1*, q_i at fair prices satisfies, for

all (\mathbf{p}, θ) ,

$$\int_T w_t \mathbf{1}[\mathbf{r}_{it} < (1, \dots, 1)] dt \leq q_i(1, \dots, 1; \mathbf{p}, \theta) \leq \int_T w_t \mathbf{1}[\mathbf{r}_{it} \leq (1, \dots, 1)] dt.$$

By *Assumption 1*, the lower bound on q_i in this expression is bounded away from zero while the upper bound is bounded below $\int_T w_t dt$. In addition, aggregate investment is bounded by the aggregate endowment in the market:

$$\sum_j q_j(ER_1, \dots, ER_n; \mathbf{p}, \theta) \leq \int_T w_t dt.$$

This implies that, the market share of security i at fair prices, which is given by

$$s_i(1, \dots, 1; \mathbf{p}, \theta) = \frac{q_i(1, \dots, 1; \mathbf{p}, \theta)}{\sum_j q_j(1, \dots, 1; \mathbf{p}, \theta)},$$

must be bounded above zero and below one. Define

$$\underline{s} := \min_i \inf \{s_i(1, \dots, 1; \mathbf{p}, \theta) : (\mathbf{p}, \theta) \in (0, 1) \times \Theta\} > 0.$$

To prove part (i) of the theorem, let $p_i > \bar{q} := 1 - \underline{s}$ and suppose the theorem does not hold, i.e., $p_i \leq \rho_i = s_i$. Then, since $\sum_{j \neq i} s_j = 1 - s_i$ and $\sum_{j \neq i} p_j = 1 - p_i$ there must be at least a security h with $ER_h \geq 1$. Pick the one with the highest expected returns. For this security we must have $\frac{ER_h}{ER_j} \geq 1$ for all j , and thus it must be that $s_h \geq \underline{s}$. But then, since $s_i \geq p_i > 1 - \underline{s}$ by assumption, we have that $s_h \leq \sum_{j \neq i} s_j < \underline{s}$, a contradiction. Obviously if security i is strictly underpriced, we must have that the remaining securities are overpriced on average: $\sum_{j \neq i} \rho_j > \sum_{j \neq i} p_j$.

Now consider the introduction of transaction costs. If, before introducing them, trader t was indifferent between investing on any security i or not investing,⁴⁸ we must have

$$\frac{ER_i}{ER_j} = \frac{r_{it}}{r_{jt}} \quad \forall j \neq i. \quad (16)$$

In this context, let $ER_i = \frac{p}{s_i^*}$, where s_i^* is the market share associated with zero transaction costs. After introducing positive transaction costs, for the agent to be indifferent between investing in security i and not investing in the market, expected returns need to satisfy

$$ER_i = (1 - \tau) \frac{p_i}{s_i'} = \frac{p}{s_i^*},$$

⁴⁸Recall from the proof of *Lemma 1* that for an agent to be indifferent between investing or not she must also be indifferent between investing among any two securities.

where s'_i is the new market share for security i , implying that $s'_i < s_i^*$. But this means that (i) the agent does not longer invest in the market at the old prices ($\rho_i = s_i^*$ for all i), and (ii) if the agent is indifferent between investing in i or not at the new prices, she has strict incentives not to invest in some of the other securities. This is because $s'_i < s_i^*$ implies that $s'_j > s_j^*$ for some security j , given that market shares add up to one. Therefore, the above indifference condition translates into

$$\frac{ER_i}{ER_j} \geq \frac{r_{it}}{r_{jt}} \quad \forall j \neq i, \quad (17)$$

with strict inequality for some j .

Summing up, the introduction of positive transaction costs can potentially reduce the demand for any given asset. If this reduction is big enough the equilibrium market share for some security may not be bounded above zero, also implying that the other securities' shares may not be bounded below one, so that the FLB may not hold. Thus, in order to show that the FLB holds in equilibrium we need to show that this does not happen for low enough τ .

The first thing to note is that, since $L_i[z|\mathbf{p}, \theta] = Pr[\mathbf{r}_{it} \ll z|\mathbf{p}, \theta]$ and $L_i[1|\mathbf{p}, \theta]$ is bounded away from zero, we can always find $z = (z_1, \dots, z_n) \ll (1, \dots, 1)$ close enough to $(1, \dots, 1)$ such that $L_i[z|\mathbf{p}, \theta]$ is also bounded above from zero for all i .

Second, notice that the smaller the track take τ , the closer the market share s'_i gets to s_i^* with $s'_i \rightarrow s_i^*$ as $\tau \rightarrow 0$ for all i , and thus the closer the indifference condition (17) gets to condition (16). Hence, if we fix $ER_i/ER_j = 1$ for all i, j , given $z \ll (1, \dots, 1)$ we can always find a low enough $\bar{\tau} > 0$ such that, for all $\tau < \bar{\tau}$, the marginal traders indifferent between investing in i or not have indifference ratios satisfying $z < \mathbf{r}_{it} < 1$.

Therefore, combining these two facts we can show that market shares when $ER_i/ER_j = 1$ for all i, j are bounded away from zero for all $\tau < \bar{\tau}$ given some small $\bar{\tau} > 0$ and apply the same reasoning as in the case of $\tau = 0$ to show that the FLB must hold in equilibrium. \square

A.1 Weakening Assumption 1

As we mention in Section 3, when endowments and beliefs are independent, the FLB could obtain even when the mass of agents strictly preferring security i over the alternatives dwindles to zero as $p_i \rightarrow 0$. The next assumption and theorem formalize this intuition.

Assumption 2. *There exist $\underline{p} > 0$ and $\alpha > 1$ such that $L_i[(1, \dots, 1) | \mathbf{p}, \theta] > \alpha p_i$ for all $p_i < \underline{p}$ and all $i = 1, \dots, n$.*

Theorem 3. *If Assumption 2 holds and endowments are independent of P , there exists $\bar{\tau} > 0$ such that for all $\tau < \bar{\tau}$ a necessary consequence of equilibrium is that there exists*

$\bar{q} < 1$ such that, for all $i = 1, \dots, n$, if $p_i > \bar{q}$ then security i is underpriced while securities $j \neq i$ are overpriced on average.

Proof. The same argument as the one used in the proof of *Theorem 1* follows through by noticing that the orthogonality of endowments and beliefs implies that s_i is bounded below by L_i . To see why, focus on the case of $\tau = 0$ and notice that

$$\begin{aligned} s_i &\geq \frac{1}{\int_T w_t dt} \int_T w_t \mathbf{1}\{r_{it}/r_{jt} < ER_i/ER_j \forall j \neq i\} dt \\ &= \frac{1}{\int_T w_t dt} \left(\int_T w_t dt \right) L_i[(\dots, ER_i/ER_j, \dots) \mid \mathbf{p}, \theta], \end{aligned}$$

where the last inequality follows from the independence of endowments and indifference ratios. But then, the above assumption guarantees that, at fair prices, $s_i \geq L_i[(1, \dots, 1) \mid \mathbf{p}, \theta] > p_i$ for all \mathbf{p} with $p_i < \underline{p}$. Given this, the same argument by contradiction used in the proof of *Theorem 1* immediately applies. \square

B Public Information and the FLB

In this section we illustrate how the FLB is mitigated by the arrival of information, which is consistent with the above findings regarding maiden and non-maiden races.

Consider the following scenario. There is going to be a sequence of two A-D security markets involving the same securities —the same result readily generalizes to having more than two markets. True payout probabilities \mathbf{p} are independent and the same across markets. We assume that the first time the market is run, agents have heterogeneous posterior beliefs. We also assume that the distribution of those beliefs is continuous with full support in $int\Delta^{n-1}$ and that $\tau = 0$. Markets are identical, *except* in the amount of public information available: before trading in the second market, all agents observe the realized outcome in the first market, i.e., which security paid positive returns. Let \mathcal{H}_i denote the event that security i pays out in the first market. Given agent t 's beliefs, let $\mathcal{L}_{ij}^t(\mathcal{H}_i)$ denote the likelihood ratio associated to outcome \mathcal{H}_i . That is, $\mathcal{L}_{ij}^t(\mathcal{H}_i) = \frac{Pr^t(\mathcal{H}_i|i)}{Pr^t(\mathcal{H}_i|j)}$, where $Pr^t(\mathcal{H}_i|j)$ is agent t 's probability assessment of observing \mathcal{H}_i conditional on security j paying out in the second market. According to Bayes' rule, the ratio of agent t 's subjective probabilities conditional on observing \mathcal{H}_i satisfies⁴⁹

$$\frac{\pi'_{it}}{\pi'_{jt}} = \mathcal{L}_{ij}^t(\mathcal{H}_i) \frac{\pi_{it}}{\pi_{jt}}, \quad i, j = 1, \dots, n, \quad (18)$$

⁴⁹We are implicitly assuming that, in the absence of public information, agents' inferences and equilibrium prices in both markets would be the same.

where π and π' represent beliefs before and after observing \mathcal{H}_i , respectively. Our next result shows that, whenever agents agree on the ‘direction’ of the updating, the release of public information about past outcomes decreases the underpricing of favorites and the average overpricing of longshots, that is, the FLB is mitigated. First, we provide the formal notions of agreement and unbiased interpretation of information we use in the result.

Definition 1. Agents *ordinally agree* on the interpretation of \mathcal{H}_i if $\mathcal{L}_{ij}^t(\mathcal{H}_i) < (>) 1$ for some t implies $\mathcal{L}_{ij}^k(\mathcal{H}_i) < (>) 1$ for all k and all i, j . Agents ordinally agree if they agree on the interpretation of \mathcal{H}_i for all $i = 1, \dots, n$.

Notice that ordinal agreement is weaker than requiring beliefs to be concordant (Milgrom and Stokey, 1982; Ottaviani and Sorensen, 2010a), which would imply ‘cardinal’ agreement, i.e. $\mathcal{L}_{ij}^t = \mathcal{L}_{ij}^k$ for all i, j, t and k .

Definition 2. Agent t is *unbiased* if $\mathcal{L}_{ij}^t(\mathcal{H}_i) > 1$ for all $j \neq i$ and all $i = 1, \dots, n$.

The notion of unbiased beliefs implies that after observing security i pay out a trader revises upwards her beliefs about i paying out in the second market relative to all other securities. In the presence of the FLB, this leads to both less underpricing of heavy favorites and less overpricing on average of longshots. That is, public information mitigates the FLB, as it is the case in *Figure 5*.

Proposition 1. *If agents are unbiased and ordinally agree then there exists \bar{p} such that for all $p_i > \bar{p}$ the expected price of security i in the second market is higher than in the first market, and the prices of securities $j \neq i$ are lower on average.*

This result implies that when *Theorem 1* holds, the release of information mitigates the FLB: heavy favorites are less underpriced and longshots exhibit less overpricing on average.

Proof. Fix the state of the world $\mathbf{p} = (p_1, p_2, \dots, p_n)$. For simplicity, we assume that, when indifferent between investing or not in the market, all agents decide to invest.⁵⁰ Given (18), market shares (i.e. prices) in the second market after observing \mathcal{H}_i are given by

$$\begin{aligned} s_i &= \frac{1}{\int_T w_t dt} \int_T w_t \mathbf{1} \left[\left(\frac{\pi'_{it}}{\pi'_{1t}}, \dots, \frac{\pi'_{it}}{\pi'_{nt}} \right) < \left(\frac{s_i}{s_1}, \dots, \frac{s_i}{s_n} \right) \right] dt \\ &= \frac{1}{\int_T w_t dt} \int_T w_t \mathbf{1} \left[\left(\frac{\pi_{it}}{\pi_{1t}}, \dots, \frac{\pi_{it}}{\pi_{nt}} \right) < \left(\mathcal{L}_{1i}^t(\mathcal{H}_i) \frac{s_i}{s_1}, \dots, \mathcal{L}_{ni}^t(\mathcal{H}_i) \frac{s_i}{s_n} \right) \right] dt. \end{aligned}$$

⁵⁰The proof logic would be the same as long as the arrival of public information does not alter agents’ decision to participate when indifferent.

It is straightforward to check that, when the distribution of prior beliefs is continuous and has full support, the release of public information leads to a higher s_i when $\mathcal{L}_{ij}^t(\mathcal{H}_i) > 1$ for all j and all t : at any given price, the mass of agents that would consider security i the optimal investment has gone up after observing the information.

Next, notice that we must have $Pr(\mathcal{H}_i) \rightarrow 1$ as $p_i \rightarrow 1$. Thus, if agents ordinally agree and are unbiased then $Pr(\mathcal{L}_{ij}^t(\mathcal{H}_i) > 1 \forall j \neq i) \rightarrow 1$ as $p_i \rightarrow 1$, implying that the probability that $\rho_i (= s_i)$ is higher in the second market is close to one for p_i close to one. That is, for p_i sufficiently high the expected price of security i is higher in the second market than in the first market. This, in turn, implies that $\sum_{j \neq i} \rho_j$ goes down in expectation. \square

C General Approach to Estimating Heterogeneity

In this section we describe how to conduct step 1 in the ML estimation of an heterogeneous population model laid out in Section 5, which involves solving the system of market share equations (7):

$$s_i^k = \frac{1}{\int_T w_t dt} \int_T w_t \mathbf{1} \left[U_t(p_i^k, R_i^k) > U_t(p_j^k, R_j^k) \forall j \neq i \right] dt, \quad i = 1, \dots, n^k.$$

Accordingly, the goal is to characterize U_t so that we can estimate the heterogeneous population model by inverting the system of equations (7) to recover \mathbf{p}^k . To do so, we assume agents have expected utility preferences and exhibit heterogeneity in beliefs and in risk attitudes. Accordingly, agent t 's payoff from investing in security i is given by

$$U_t(p_i, R_i) = \pi_{it} u(w_t R_i, \gamma_t) + (1 - \pi_{it}) u(-w_t, \gamma_t), \quad (19)$$

where π_{it} is agent t 's belief about security i and u is utility over wealth with risk attitudes governed by the one-dimensional parameter γ_t —e.g. the coefficient of absolute risk aversion in CARA utility. The cardinality of expected utility allows to normalize the utility from losing w_t to be zero, i.e., $u(-w_t, \gamma_t) = 0$. In addition, since agent preferences over securities in a given market are invariant to a monotonic transformation of utility, we can take the log of the RHS of (19) and subtract $\log p_1$ to write U_t as

$$U_t(p_i, R_i) = \log u(w_t R_i, \gamma_t) + \xi_i + \nu_{it},$$

where $\xi_i = \log p_i - \log p_1$ and $\nu_{it} = -\log r_{it}$. If we assume that risk parameters and beliefs are independently distributed and that agents have equal endowments (normalized to 1),⁵¹

⁵¹Alternatively, we could assume $w_t = a > 0$ for all $t \in T$. Generally, without information on endowments, the risk parameter is not separately identified from a .

we can write (7) as

$$s_i^k = \int_{\gamma_t \in \mathbb{R}} \int_{\nu_t \in \mathbb{R}^n} \mathbf{1} \left[\log u(R_i^k, \gamma_t) + \xi_i + \nu_{it} > \log u(R_j^k, \gamma_t) + \xi_j + \nu_{jt} \forall j \neq i \right] dP(\nu_t; \theta) dH(\gamma_t),$$

for all $i = 1, \dots, n^k$, where $P(\cdot; \theta)$ is the (continuous) distribution of $\nu_t = (\nu_{1t}, \dots, \nu_{n^k t})$ and H represents the distribution of risk parameter γ_t . Thus, as long as this system of market shares has a unique solution in terms of the vector $(\xi_1, \dots, \xi_{n^k})$, we can recover the underlying probabilities \mathbf{p}^k from observed odds \mathbf{R}^k .

One way to ensure the existence of a unique solution is to assume, as we do in our estimation, that ν_t is a variance mixture of independent logit errors with mixing distribution $G(\sigma)$. Accordingly, market shares are given by

$$s_i^k = \int_{\gamma_t \in \mathbb{R}} \int_{\sigma \in \mathbb{R}_+} \frac{\exp\left(\frac{1}{\sigma} \log u(R_i^k, \gamma_t) + \frac{1}{\sigma} \xi_i\right)}{\sum_{j=1}^{n^k} \exp\left(\frac{1}{\sigma} \log u(R_j^k, \gamma_t) + \frac{1}{\sigma} \xi_j\right)} dG(\sigma) dH(\gamma), \quad i = 1, \dots, n^k. \quad (20)$$

This system has a solution, given that it satisfies the sufficient conditions in Berry (1994), and it is unique as shown by Berry et al. (2011).

D Non-parametric Identification of the Belief Heterogeneity Model

In this section we formally show how the identification of the distribution of beliefs, which is the key primitive we estimate in our model, is obtained from the market level data we observe. The key source for identification, which our empirical strategy for estimating heterogeneous beliefs directly exploits, stems from variation in the fundamentals $\mathbf{p} = (p_1, \dots, p_n) \in \text{int}\Delta^{n-1}$, i.e., the state probabilities of the underlying Arrow-Debreu securities, across different markets. The present appendix contributes to the paper by making precise how variation in the fundamentals \mathbf{p} *across* markets can identify the distribution of beliefs that operates *within* markets.

In order to focus our attention on the logic of identification, we consider here a simplified setting with two horses and track take $\tau = 0$. Our discussion however easily generalizes to the n -horse context. Indeed, the additional variation made possible by n -horses rather than two horses only aids identification rather than complicating it (observe, for example, that an n -horse can always replicate a two horse race by letting $n-1$ horses have arbitrarily similar state probabilities but not vice versa) and thus the two horse problem is the essential setting to study.

In each race (i.e. market) we have two horses labeled $i = 1, 2$ and can in principle observe the market share s of horse 1 (the share of horse 2 being simply $s_2 = 1 - s$) as well as the vector of expected returns (ER_1, ER_2) .⁵² Letting $\delta = \log ER_1 - \log ER_2$, the model of belief heterogeneity we estimate in the paper relates these market observables via the random utility model

$$u_i = \delta + \epsilon_i$$

where ϵ_i corresponds to the belief disturbance associated with agent i . Agent i in the market bets on horse 1 if $u_i > 0$. Assuming a continuum of agents in the market and ϵ_i is distributed independently of (ER_1, ER_2) with distribution function H , we thus have that

$$s = H(\delta). \tag{21}$$

If δ varies across markets such that its support covers all of \mathbb{R} , i.e., for each $c \in \mathbb{R}$ there exists a market (s^c, δ^c) such that $\delta^c = c$, then it is clear that H is non-parametrically identified. In particular for any $c \in \mathbb{R}$, we have that $s^c = H(c)$. Thus δ plays the role of the well known “special regressor” that is a key device used in the literature to non-parametrically identify discrete choice models (see e.g., Lewbel (2000); Matzkin (2007)).

However the variation in δ in the data is constrained by the equilibrium of the model and thus δ differs from the standard “special regressor”, which is assumed to vary freely without such an equilibrium constraint. In our case, the primitive that varies exogenously across markets is $p \in (0, 1)$, which is the probability of horse 1 winning (where $p_2 = 1 - p$ is thus the probability of horse 2 winning). Recalling that $ER_i = \frac{p_i}{s_i}$, we have that equilibrium prices and quantities are determined by the solution to the fixed point condition

$$s = H\left(\log \frac{p}{s} - \log \frac{1-p}{1-s}\right). \tag{22}$$

Assuming that $H \in \mathbb{H}$, where \mathbb{H} denote the set of distribution functions that are continuous and have full support on \mathbb{R} , then it is straightforward to see that a unique solution to (22) exists with $s \in (0, 1)$ for a given $p \in (0, 1)$.⁵³ Thus the equilibrium manifold can be expressed as a function $s(p; H)$ for $p \in (0, 1)$, where the manifold as a whole depends on the underlying distribution of beliefs H . We now show that for any distribution function $H \in \mathbb{H}$, the equilibrium manifold it generates induces variation in $\delta(p; \delta)$ that satisfies the full support requirement of a special regressor. That is, for any $H \in \mathbb{H}$, it is indeed the

⁵²Recall the discussion at the start of Section 5 concerning how returns are identified in the data in addition to market shares.

⁵³Observe that the LHS is continuous and strictly increasing in s while the RHS is continuous and decreasing. Furthermore, at $s = 0$ (or $s = 1$), the LHS is 0 (or 1) while the RHS is 1 (or 0). Thus, the LHS and the RHS intersect exactly once, with the fixed point lying in $(0,1)$.

case that $\delta(p; H)$ sweeps over \mathbb{R} as p sweeps over $(0, 1)$. As a consequence, for any $H \in \mathbb{H}$, H can be recovered from the manifold it induces through the relationship (21).

The key lemma we use to establish this result is the following.

Lemma 2. *For any $H \in \mathbb{H}$, the set $\{\delta(p; H) \mid p \in (0, 1)\}$ is not bounded from above or below.*

Proof. For simplicity let us suppress the dependence of the manifold on H as it is being held fixed throughout the argument. Assume there exists $\bar{\delta}$ such that $\delta(p) \leq \bar{\delta}$ for all $p \in (0, 1)$. Then for all $p \in (0, 1)$ we have that $s_2(p) = 1 - s(p) \geq 1 - H(\bar{\delta}) > 0$. Hence $\frac{p_2}{s_2(p)} \rightarrow 0$, or likewise $\log ER_2(p) \rightarrow -\infty$ as $p \rightarrow 1$ (because $p_2 = 1 - p$ goes to zero while s_2 stays bounded away from zero). Moreover because shares sum to 1 we must have $\log ER_1(p) > 0$ for all p sufficiently close to 1. Hence $\delta(p) \rightarrow \infty$ as $p \rightarrow 1$, which contradicts the presence of an upper bound. The proof of the non-existence of a lower bound follows similarly. \square

Hence $\delta(p)$ is not bounded above or below as $p \in (0, 1)$. Furthermore, $\delta(p)$ is continuous as it is the outcome of the unique fixed point of (22) where the fixed point problem varies continuously in p . Thus, by the intermediate value theorem, for any $c \in \mathbb{R}$ there exists a p^c such that $\delta(p^c) = c$. Hence $\delta(p)$ has full support over \mathbb{R} and H is identified non-parametrically.

In the model we estimate in the paper, we restrict H further to be a variance mixture of logistic distributions. That is for any $t \in \mathbb{R}$ we model

$$H(t) = \int_{\sigma > 0} F(t \mid \sigma) dG(\sigma).$$

where $F(\cdot \mid \sigma)$ is a standard logistic distribution with scale parameter σ . The object we estimate is G , i.e., the distribution of “types” in the population where each type corresponds to a certain level of variance in beliefs, where G is assumed to have finite support. Let \mathbb{G} denote the set of all distribution functions with finite support (which is a non-parametric class of distributions as it is infinite dimensional). Using standard results from the identifiability of mixture models (see e.g., Shi et al. (2012)), we can recover our structural object of interest G from H .

E Closed Form Identification of the HR Model

In this section we describe the steps needed to back out payout probabilities $\mathbf{p} = (p_1, \dots, p_n)$ from observed odds $\mathbf{R} = (R_1, \dots, R_n)$. Two steps are needed. In the first one we compute

the utility parameters of the marginal bettors $\{\gamma_i\}_{i=1}^n$ using the market shares associated to \mathbf{R} . In the second step we use the indifference conditions of marginal bettors to recover the probabilities.

Step 1: Identifying Marginal Bettors Using (2) we obtain the market shares $\mathbf{s} = (s_1, \dots, s_n)$. Recall that market shares satisfy (12), i.e.,

$$\begin{aligned} s_1 &= F(\gamma_1; \theta) \\ s_2 &= F(\gamma_2; \theta) - F(\gamma_1; \theta) \\ &\vdots \\ s_n &= 1 - F(\gamma_{n-1}; \theta). \end{aligned}$$

Since $F(\cdot|\theta)$ is strictly increasing in its support we can invert this system of equations and obtain the utility parameters:

$$\begin{aligned} \gamma_1 &= F^{-1}(s_1; \theta) \\ \gamma_2 &= F^{-1}(s_1 + s_2; \theta) \\ &\vdots \\ \gamma_n &= F^{-1}(s_1 + \dots + s_{n-1}; \theta). \end{aligned}$$

Step 2: Recovering the Probabilities From (11) we know that p_i satisfies

$$\begin{aligned} U(p_i, R_i, \gamma_i) &= U(p_{i+1}, R_{i+1}, \gamma_i) \\ \Rightarrow p_i u(R_i, \gamma_i) + (1 - p_i)u(-1, \gamma_i) &= p_{i+1}u(R_{i+1}, \gamma_i) + (1 - p_{i+1})u(-1, \gamma_i) \\ \Rightarrow p_{i+1} &= \left(\frac{u(R_i, \gamma_i) - u(-1, \gamma_i)}{u(R_{i+1}, \gamma_i) - u(-1, \gamma_i)} \right) p_i. \end{aligned}$$

For $i = 1, \dots, (n - 1)$, define

$$c_i = \prod_{j=1}^i \frac{u(R_j, \gamma_j) - u(-1, \gamma_j)}{u(R_{j+1}, \gamma_j) - u(-1, \gamma_j)}.$$

That is, we have that $p_{i+1} = c_i p_i$. Thus, if we recover p_1 we can compute the remaining probabilities recursively. Since probabilities add up to one, we must have that

$$p_1 + c_1 p_1 + \dots + c_{n-1} p_1 = 1 \quad \Rightarrow \quad p_1 = \frac{1}{1 + c_1 + \dots + c_{n-1}}.$$

Thus for a given family of expected utility preferences $U(\cdot, \cdot, \gamma)$, and a given cdf $F(\cdot; \theta)$ over γ , and a given market with $R_1 > \dots > R_n$, we can analytically solve for the underlying state of nature $(p_1, \dots, p_n) \in \text{int}\Delta^{n-1}$.

References

- Ali, Mukhtar**, “Probability and Utility Estimates for Racetrack Betting,” *Journal of Political Economy*, 1977, 85 (4), 803–815.
- Aliprantis, C.D. and K.C. Border**, *Infinite dimensional analysis: a hitchhiker’s guide*, Springer Verlag, 2006.
- Berry, Steven T.**, “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 1994, pp. 242–262.
- , **Gandhi Amit K., and Philip A. Haile**, “Connected Substitutes and Invertibility of Demand,” *Working paper*, 2011.
- , **James Levinsohn, and Ariel Pakes**, “Automobile Prices in Market Equilibrium,” *Econometrica*, July 1995, 63 (4), 841–90.
- Blackwell, David and Lester Dubins**, “Merging of Opinions with Increasing Information,” *Annals of Mathematical Statistics*, September 1962, 33 (3), 882–886.
- Bombardini, Mathilde and Francesco Trebbi**, “Risk Aversion and Expected Utility Theory: An Experiment with Large and Small Stakes,” *Journal of the European Economic Association*, Forthcoming.
- Camerer, Colin F.**, “Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting,” *Journal of Political Economy*, 1998, 106 (3), 457–482.
- Cameron, A.C. and P.K. Trivedi**, *Microeconometrics: methods and applications*, Cambridge Univ Pr, 2005.
- Chiappori, Pierre-Andre, Amit Gandhi, Bernard Salanie, and Francois Salanie**, “Identifying Preferences under Risk from Discrete Choices,” *American Economic Review: Papers & Proceedings*, 2009, 99 (2), 356–362.
- Cipriani, Marco and Antonio Guarino**, “Noise Trading in A Laboratory Financial Market: A Maximum Likelihood Approach,” *Journal of the European Economic Association*, May 2005, 3 (2-3), 315–321.

- Cochrane, John H.**, “Efficient Markets Today,” 2007. Conference on Chicago Economics.
- Cowgill, Bo, Justin Wolfers, and Eric Zitzewitz**, “Using Prediction Markets to Track Information Flows: Evidence from Google,” 2009. Mimeo.
- Easley, David and Lawrence Blume**, “Learning to be Rational,” *Journal of Economic Theory*, April 1982, *26* (2), 340–351.
- , **Nicholas M. Kiefer, and Maureen O’hara**, “One Day in the Life of a Very Common Stock,” *Review of Financial Studies*, 1997, *10* (3), 805–835.
- Feldman, Mark**, “An Example of Convergence to Rational Expectations with Heterogeneous Beliefs,” *International Economic Review*, 1987, *28* (3), pp. 635–650.
- Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright**, “Anatomy of an Experimental Political Stock Market,” *American Economic Review*, December 1992, *82* (5), 1142–1161.
- Glosten, Lawrence R. and Paul R. Milgrom**, “Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders,” *Journal of Financial Economics*, 1985, *14*, 71–100.
- Golec, Joseph and Maury Tamarkin**, “Bettors Love Skewness, Not Risk, at the Horse Track,” *Journal of Political Economy*, 1998, *106* (1), 205–225.
- Gossner, Olivier and Tristan Tomala**, “Entropy bounds on Bayesian learning,” *Journal of Mathematical Economics*, 2008, *44* (1), 24 – 32.
- Griffith, Richard M.**, “Odds Adjustment by American Horse-Race Bettors,” *American Journal of Psychology*, 1949, *62* (2), 290–294.
- Jullien, Bruno and Bernard Salanié**, “Estimating Preferences under Risk: The Case of Racetrack Bettors,” *Journal of Political Economy*, 2000, *108* (3), 503–530.
- and —, “Empirical Evidence on the Preferences of Racetrack Bettors,” in D.B. Hausch and W.T. Ziemba, eds., *Handbook of sports and lottery markets*, Elsevier/North-Holland, 2008.
- Kalai, Ehud and Ehud Lehrer**, “Weak and strong merging of opinions,” *Journal of Mathematical Economics*, 1994, *23* (1), 73 – 86.
- Keim, Donald B.**, “financial market anomalies,” in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, Basingstoke: Palgrave Macmillan, 2008.

- Kyle, Albert S.**, “Continuous Auctions and Insider Trading,” *Econometrica*, November 1985, *53* (6), 1315–1336.
- Lehrer, Ehud and Rann Smorodinsky**, “Merging and Learning,” *Lecture Notes-Monograph Series*, 1996, *30*, pp. 147–168.
- Lewbel, Arthur**, “Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables,” *Journal of Econometrics*, 2000, *97* (1), 145–177.
- Long, J. Bradford De, Andrei Shleifer, Lawrence H. Summers, and Robert T. Waldmann**, “Noise Trader in Financial Markets,” *Journal of Political Economy*, August 1990, *98* (4), 703–738.
- Manski, Charles F.**, “Measuring expectations,” *Econometrica*, 2004, *72* (5), 1329–1376.
- Matzkin, Rosa L.**, “Heterogeneous choice,” *Econometric Society Monographs*, 2007, *43*, 75.
- Milgrom, Paul R. and Nancy Stokey**, “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, 1982, *26*, 17–27.
- Mitchell, Dick**, *Winning Thoroughbred Strategies*, W. Morrow (New York), 1989.
- Ottaviani, Marco and Peter N. Sorensen**, “Parimutuel versus Fixed-Odds Markets,” April 2005. Mimeo.
- and – , “The Timing of Parimutuel Bets,” 2006. Mimeo.
- and – , “Aggregation of Information and Beliefs: Asset Pricing Lessons from Prediction Markets,” 2010. Mimeo.
- and – , “Noise, Information, and the Favorite-Longshot Bias in Parimutuel Predictions,” *American Economic Journal: Microeconomics*, February 2010, *2* (1), 58–85.
- Potters, Jan and Jörgen Wit**, “Bets and Bids: Favorite-Longshot Bias and Winner’s Curse,” 1996. Mimeo.
- Quandt, Richard E.**, “Betting and Equilibrium,” *Quarterly Journal of Economics*, 1986, *101* (1), 201–207.
- Sadka, Ronnie and Anna Scherbina**, “Analyst Disagreement, Mispricing, and Liquidity,” *The Journal of Finance*, 2007, *62* (5), 2367–2403.

- Sandroni, Alvaro and Rann Smorodinsky**, “The Speed of Rational Learning,” *International Journal of Game Theory*, 1999, 28 (2), 199–200.
- Sauer, Raymond D.**, “The Economics of Wagering Markets,” *Journal of Economic Literature*, December 1998, 36 (4), 2021–2064.
- Serrano-Padial, Ricardo**, “Naive Traders and Mispricing in Prediction Markets,” *Journal of Economic Theory*, 2012, *Forthcoming*.
- Shi, Xiaoxia**, “A Non-Degenerate Vuong Test,” 2012. Mimeo.
- Shi, ZiQiang, TieRan Zheng, and JiQing Han**, “Identifiability of multivariate logistic mixture models,” *ArXiv e-prints*, August 2012.
- Shin, Hyun Song**, “Optimal Betting Odds Against Insider Traders,” *Economic Journal*, 1991, 101 (408), 1179–1185.
- , “Prices of State Contingent Claims with Insider Traders and the Favourite Longshot Bias,” *Economic Journal*, 1992, 102 (411), 426–435.
- Snowberg, Erik and Justin Wolfers**, “Explaining the Favorite-Longshot Bias: Is it Risk-Love, or Misperceptions?,” *Journal of Political Economy*, August 2010, 118 (4), 723–746.
- Stein, Jeremy C.**, “Presidential Address: Sophisticated Investors and Market Efficiency,” *The Journal of Finance*, 2009, 64 (4), 1517–1548.
- Thaler, Richard H. and William T. Ziemba**, “Anomalies: Parimutuel betting markets: Racetracks and lotteries,” *Journal of Economic Perspectives*, 1988, 2 (2), 161–174.
- Tompkins, Robert G., William T. Ziemba, and Stewart D. Hodges**, “The Favorite-Longshot Bias in S&P 500 and FTSE 100 Index Futures Options: The Return to Bets and the Cost of Insurance,” in Donald B. Hausch, ed., *Handbook of Sports and Lottery Markets*, Elsevier B. V., 2008, chapter 9, pp. 161–178.
- Townsend, Robert M.**, “Market Anticipations, Rational Expectations, and Bayesian Analysis,” *International Economic Review*, 1978, 19 (2), pp. 481–494.
- Train, Kenneth E.**, *Discrete choice methods with simulation*, Cambridge University Press, 2003.
- Vives, Xavier**, “How Fast do Rational Agents Learn?,” *The Review of Economic Studies*, 1993, 60 (2), pp. 329–347.

Vuong, Quang H., “Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses,” *Econometrica*, 1989, *57* (2), pp. 307–333.

White, Halbert, “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica*, 1982, *50* (1), pp. 1–25.

– , “Specification testing in dynamic models,” in “Advances in econometrics, Fifth world congress,” Vol. 1 Cambridge, Cambridge Press 1987, pp. 1–58.