Individual Behavior and Beliefs in Experimental Parimutuel Betting Markets*

Frédéric Koessler, Charles Noussair, and Anthony Ziegelmeyer†

Abstract

We study experimental parimutuel betting markets with asymmetrically informed bettors. We propose a theoretical model, the *Adaptive Model*, which serves as our source of null hypotheses about individual behavior and the capacity of the markets to aggregate information. In one treatment, groups of eight participants bet against each other in twenty repetitions of a sequential betting market. The second treatment is identical, except that bets are observed by other participants who assess the winning probabilities of each outcome. In the third treatment, the same individuals place bets and assess the winning probabilities of the outcomes. A favorite-longshot bias is observed in the first and second treatments, but it is sharply reduced in the third treatment. Information aggregation is better in the third than in the other two treatments, because contrarian betting is almost completely eliminated by the belief elicitation procedure. Placing bets improves the accuracy of belief statements. A statistical generalization of the Adaptive Model explains the data very effectively.

Keywords: belief elicitation, information aggregation, parimutuel betting, experiment.

JEL Classification: C90, C91, C92, D40, D8, D82.

*The former title of this paper was "Information Aggregation in Parimutuel Betting Markets: An Experimental Study using Elicited Beliefs". We thank participants at seminars in Bonn, Cergy-Pontoise, Jena, Paris, at the GEW-Tagung in Wittenberg, at the AFFI International Conference 2004, and at the ESA meetings in Boston, Strasbourg and Erfurt for very useful comments. Most of all we thank Stéphane Bertrand and Kene Boun My for research assistance.

†Koessler: THEMA (CNRS), Université de Cergy-Pontoise, 33 Boulevard du Port, F-95011 Cergy-Pontoise, France (email: Frederic.Koessler@eco.u-cergy.fr). Noussair: Department of Economics, Emory University, Atlanta, GA 30322-2240, USA (email: cnoussa@emory.edu). Ziegelmeyer: Max Planck Institute for Research into Economic Systems, Strategic Interaction Group, Jena, Germany (email: ziegelmeyer@mpiew-jena.mpg.de).

1 Introduction

As the principal means of wagering on racetracks throughout the world, parimutuel betting markets are of considerable empirical importance. In the United States alone, the value of bets on horse races in 2003 was 17.83 billion dollars (International Federation of Horseracing Authorities, http://www.horseracing-intfed.com).[1] However, parimutuel betting markets are of interest in the economics and finance research communities primarily because they can be viewed as simple representations of financial markets (Thaler and Ziemba, 1988). Both financial and parimutuel betting markets are characterized by investors (bettors) with potential access to public, and possibly private, information, and by uncertainty about future payoffs (although betting markets have the property that there is a known termination point in which payoffs become known with certainty). As a result, a bettor faces a rich decision problem that shares the essential features of the one that confronts an investor in a financial market: is it better for an individual to be a net purchaser or a net seller given (a) the current market price, (b) his own private information, (c) his beliefs about others' degree of rationality, and (d) the history of trading activity?

The experiment reported in this paper focuses on the influence of the elements (a) - (d) on both individual behavior and the resulting market outcomes in a parimutuel betting market. The setting is a straightforward extension of experimental designs used to study information cascades (Anderson and Holt, 1997; Hung and Plott, 2001), with the additional feature that market prices, the odds on each outcome, exist and are updated after each individual makes his bet. We consider three different experimental treatments in which individuals observe private signals that reveal information about which of two a priori equally likely events has occurred. In the first treatment, individuals sequentially place bets on the two possible outcomes, and a winning bet yields a return that is decreasing with the proportion of individuals who have bet on the same outcome. In the second treatment, there are two groups of individuals: bettors and observers. Bettors participate in an experimental market similar to the one in the first treatment, while knowing that observers exist who are asked about the likelihood of each outcome based on the betting data they have observed. In the third treatment, individuals both place bets and state their beliefs about the likelihood of each outcome. A set of null hypotheses about betting behavior and beliefs for our experiment is derived from a limited foresight model, called the *Adaptive Model*, in which bettors correctly update their beliefs, based on their private information and the prior decisions of others, but act as if the odds after their choices are the final odds. Though it departs only slightly from standard game-theoretic solution concepts, the Adaptive Model is appealing as a descriptive model because, unlike Nash or sequential equilibrium, it generically makes a unique prediction at every decision node, and bettors make inferences only from past history.

The experiment considers two sets of issues. The first set of issues concerns the ability of the Adaptive Model to predict betting decisions, resulting market outcomes, and belief statements. We compare the model's predictions to those of a plausible alternative betting rule, the *Private Information Plus Odds* heuristic, in which bettors do not infer any information from past bets but simply respond optimally to the current odds given their private information. We also consider the ability of several statistical generalizations of the Adaptive Model to explain the inaccuracies in the model's predictions. The second set of issues consists of several questions regarding the impact of betting and belief elicitation on each other. First, does the elicitation of beliefs affect betting decisions and the information aggregation properties of the market, and if so, is it possible to offer a behavioral model that accounts for the effect? We address this question by comparing betting behavior and market outcomes between treatments 1 and 3, and exploring whether statistical generalizations of the Adaptive Model can account for the differences. Second, does the placing of bets improve the accuracy of the belief statements and what impact does it

[1] Moreover, markets with a parimutuel structure are used as a forecasting tool for macroeconomic variables. Since late 2002, Deutsche Bank and Goldman Sachs have been conducting Parimutuel Derivatives Call Auctions. These are markets in which individuals may bet on the realization of macroeconomic variables. Payouts depend on odds on each outcome that are updated as new bets are received. Contingent claims markets, which can be organized as parimutuel markets, have been enthusiastically proposed as a forecasting tool for firms (e.g., Time magazine, 7/6/2004).

have on the perception of others' degree of rationality? To answer this question we compare the belief statements in treatments 2 and 3 and we measure their consistency with both the market estimates and the beliefs generated from the Adaptive Model. Third, we compare betting data in treatments 1 and 2, to investigate whether the quality of betting decisions improves when bettors are observed by another group of individuals, whose payoffs depend on their betting decisions.

Our experiment complements the numerous empirical studies which use racetrack betting data to study models of market pricing and behavior under uncertainty. One line of this empirical research examines the relationship between the objective probability of winning a bet and its market estimate (see Griffith, 1949, or Jullien and Salanié, 2000), and concludes that market odds are reasonably good estimates of winning probabilities, a robust anomalous regularity called the *favorite-longshot bias* notwithstanding. The favorite-longshot bias refers to the observation that prices reflect a tendency to underbet on favorites and to overbet on longshots. The main drawback of using econometric methods on racetrack data is that individual data on bets and on bettors' characteristics are not available. This leads researchers to focus on the behavior of a "representative" bettor that captures the average risk attitude of the participants. The second line of this empirical research using racetrack betting data reveals the limitations of this model by investigating market-level informational efficiency[2] and by providing evidence that informed bettors earn higher returns (Schnytzer and Shilony, 1995). Indeed, in his survey of the empirical evidence on the properties of prices in wagering markets, Sauer (1998) argues that the empirical regularities that are inconsistent with generic notions of efficiency present a challenge to equilibrium models of the wagering market and that future theoretical models should include heterogeneous agents and information asymmetries.[3]

Field studies of parimutuel betting can test various forms of the efficient market hypothesis but they cannot determine whether or not markets successfully aggregate information, because there are no data on the private information available to the bettors. Plott, Witt, and Yang (2003), who investigate the ability of experimental parimutuel betting markets to aggregate information, emphasize that experimental studies can overcome the problem of the unobservability of private information. Their emphasis is on the identification of principles of behavior at the market rather than at the individual level, and their markets are characterized by a large number of transactions in continuous time. They find that the market aggregates information effectively in simple environments, but that it is less effective as the environment becomes more complex. Our experiment complements Plott, Witt, and Yang (2003) by focusing on detailed theories of individual behavior. Cipriani and Guarino (2005), and Drehmann, Oechssler, and Roider (2004a), also study markets with similarities to ours. They conduct experimental tests of the theory of informational cascades in financial markets, where prices are efficiently set by a market maker according to the order flow. Their findings are discussed in the conclusion.

Our results can be summarized as follows. First, pricing patterns reflect greater information aggregation in treatment 3, in which bettors must state beliefs, than in the other two treatments. Treatment 3 is characterized by less frequent failure to herd when herding is appropriate, and by less contrarian behavior. Because failure to herd and contrarian behavior involve betting on the longshot when it is suboptimal to do so, treatment 3 exhibits a much smaller favorite-longshot bias than the other two treatments. It appears that the attractive payouts for lonsghots in the event that they win, coupled with a poor estimate of the probability the longshot is victorious, leads to decision errors. However, the elicitation of beliefs improves bettors' probability assessments, which in turn discourages an overemphasis on the attractive return associated with a longshot during the decision process.[4]

[2]A market is considered informationally efficient if the price of a financial asset is an unbiased estimate of its fundamental value, reflecting all the information available to the market. Empirical tests of weak efficiency equate available information with historical prices and returns, tests of semistrong efficiency add public announcements to the set of available information, and tests of strong efficiency additionally include private information. Vaughan Williams (1999) surveys the literature addressing the issue of information efficiency in betting markets.

[3]See Ottaviani and Sorensen (2004, 2005) for recent attempts in this direction.

[4]The literature that addresses the impact of belief elicitation is only concerned with strategic uncertainty, by measuring

Our second major finding is that the Adaptive Model describes behavior very well in treatment 3, while a statistical generalization of the model, called $QRAM$, in which commonly known deviations from the Adaptive Model become less likely the costlier they are, provides a very good fit to the data from all three treatments. Limiting the steps of bettors' reasoning does not improve the model's fit, and thus decisions appear to reflect a considerable depth of reasoning. Third, belief statements are both more accurate and more consistent with market activity when the agents stating beliefs are also placing bets. In contrast, observers' beliefs in treatment 2 are consistent with more noise in decision making of bettors than the market actually exhibits. This may be the case because it is more difficult for an individual to formulate good belief estimates when he does not participate in the betting process. However, another factor may be that observers underestimate the rationality of other agents, in this context those who are placing bets. Fourth, betting decisions are no better when they are observed than when they are not observed, indicating that public knowledge of the fact that beliefs are being elicited is insufficient to improve decision-making. Bettors must be stating their own beliefs to induce an improvement in their betting decisions.

The paper is structured as follows. In the next section we describe the experimental design and the practical procedures of the experiment. In section 3 we present the Adaptive Model and the measures used in the data analysis. The results of the experiment are presented in section 4. Section 5 concludes.

2 The Experiment

2.1 The betting market

Consider a horse race with two horses called A and B. There is a finite set $N \equiv \{1, \ldots, n\}$ of bettors. Each bettor is endowed with one unit of money which he is required to wager on one of the horses. In other words, each bettor $i \in N$ chooses $s_i \in \{A, B\}$ where A or B consists of betting one unit of money on horse A or horse B, respectively. Bets are made sequentially with bettor i denoting the ith bettor in the sequence. Each bettor observes all previous bets before making his choice. Bettors are not permitted to cancel their bets after they are made. Each bettor has a flat common prior belief on the payoff-relevant state space $\{\theta_A, \theta_B\}$, where θ_A stands for "horse A wins", and θ_B stands for "horse B wins".

For any profile of bets $s = (s_1, \ldots, s_n) \in \{A, B\}^n$ and any horse $H \in \{A, B\}$, let $h(s) = |\{i \in N : s_i = H\}|$ be the number of bettors who bet on horse H, and let $\overline{H}$ denote the horse other than H. The *odds against horse* H, which is given by the total number of bets on horse $\overline{H}$ divided by the total number of bets on horse H, is denoted by

$$O_H(s) = \frac{n - h(s)}{h(s)}.$$

If bettor i bets on the winning horse, then his payoff equals the *return* of this horse, which is equal to the odds against it plus 1 (the amount bet is also returned to the bettor in addition to his winnings, and is therefore included in his payoff). If he bets on the losing horse, bettor i receives 0 payoff, losing his stake.

expectations about others' strategies in complete information games. Most of the experimental evidence suggests that prompting subjects for beliefs about others' strategies moves their choices closer to equilibrium. Warglien, Devetag, and Legrenzi (1998) and Croson (2000) show that asking subjects to guess their opponents' actions increases subjects' tendency to play according to iteratively undominated strategies. More recently, Rutström and Wilcox (2004) investigate whether belief elicitation changes the way subjects play a 2x2 game with a unique equilibrium in mixed strategies. They observe that subjects behave as if they construct mental models of their opponents that are both more sophisticated and robust in the presence of belief elicitation. Additionally, the two authors show that game-theoretically motivated learning models fit the data much better when the action data is generated jointly with belief elicitation than without. Less conclusive evidence is reported in Wilcox and Feltovich (2001) and Costa-Gomes and Weizsäcker (2005), where the belief elicitation procedure has no strong impact on subjects' behavior.

Before making his decision, each bettor i receives a *private signal* $q_i \in \{q^A, q^B\}$ that is correlated with the true state of nature. Conditional on the state of nature, bettors' signals are independent, identically distributed, and satisfy

$$\Pr(q_i = q^A \mid \theta_A, q_j) = \Pr(q_i = q^A \mid \theta_A) = \pi \in (1/2, 1)$$
$$\Pr(q_i = q^A \mid \theta_B, q_j) = \Pr(q_i = q^A \mid \theta_B) = 1 - \pi,$$

for all i, $j \in N$, $i \neq j$. Hence, once bettor i has received a signal q_i, his beliefs about the states of nature are given by $\Pr(\theta_H \mid q^H) = \pi$ and $\Pr(\theta_H \mid q^{\overline{H}}) = (1 - \pi)$, $H = A, B$.

2.2 Procedures

The experiment was conducted in eight sessions at the laboratory for experimental economics (LEES) at Louis Pasteur University in Strasbourg, France. 176 subjects, who had no previous experience with economic experiments on betting, were recruited for participation in the study. All sessions were conducted in French. Table 1 contains the number of sessions, groups, participants, bets placed and belief statements in each of the three treatments.

Table 1: Number of Sessions, Groups, Subjects, Bets, and Belief Statements in Each Treatment

	Treatment 1: Subjects only place bets (20 rounds)	Treatment 2: 1/2 of the subjects place bets, 1/2 of the subjects state beliefs (20 rounds)	Treatment 3: Subjects both place bets and state beliefs (20 rounds)
Sessions	2	4	2
Groups	4	8	4
Subjects	32	64	32
Bets	$32 \times 20 = 640$	$32 \times 20 = 640$	$32 \times 20 = 640$
Belief statements	0	$32 \times 8 \times 20 = 5120$	$32 \times 8 \times 20 = 5120$

At the beginning of each session, each subject was randomly assigned to one of two groups of eight. Group assignments remained the same for the entire session. In each session, both groups participating in treatments 1 and 3, and one of the two groups in treatment 2, bet against each other in 20 repetitions of the market described in the previous section, with $n = 8$ betting periods per repetition. We use the term *round* to refer to a repetition of the sequential betting market, while each *period* refers to one subject's turn to bet within a round. Subjects were instructed on the rules of the market and the use of the computer program[5] with written instructions. These were read aloud by an assistant, a subject who was chosen at random at the beginning of the session and did not participate in the betting or belief elicitation process, but was instead paid the average of the participants' earnings. A short questionnaire and one dry run followed.[6] Afterwards, the twenty rounds of the sequential betting market that constituted the experiment took place.[7] Communication between the subjects was not allowed. Each session was between 90 and 135 minutes in duration.

[5]The program is based on an application developed by Boun My (2002) designed for Visual Basic.

[6]During the dry run, which did not count toward subjects' earnings, subjects were approached one at a time by the assistant and asked to make choices consistent with a predetermined sequence of bets. The predetermined sequence of bets was announced publicly before the start of the dry run and the assistant stressed its random nature. The dry run was intended to ensure subjects' familiarity with the software without enabling them to make inferences about others' behavior.

[7]Twenty-two subjects were recruited for each session, and sixteen were retained for participation in the remainder of the session on the basis of their performance on the questionnaire.

2.2.1 The betting process

Each round involved the following sequence of events. At the beginning of a round, the assistant randomly chose between color A (state θ_A) and color B (state θ_B) by tossing a coin. Subjects were not made aware of the color that was chosen until the end of the round. Different colors were used in each round to be assigned to the two identifying letters, in order to reduce the likelihood that subjects believed that a dependence existed across rounds.[8] Subjects were then chosen in random order, which differed from round to round, to bet one Experimental Currency Unit (ECU) on either color A or B. Before making his bet, each subject observed the current returns and received a private signal correlated with the correct color. Private signals were not made public at any time. The probability that any signal was correct was equal to $\pi = 3/4$ for all bettors (this probability was public information), and each bettor's signal was drawn independently. On subjects' computer screens, the signal took the form of a ball drawn from an urn containing 4 balls, three "correct" balls and one "incorrect" ball. At the end of each round, the final odds, the winning color, the payout from a bet on each color, and the subject's own earnings were displayed on his computer screen.

For a given round, the same private signals were used in all sessions, although private signals differed between rounds within a given session. Thus, in a given round t, the same eight private signals conditional on the state of nature were used in each session, though the signals in round $s \neq t$ differed from those in t. At the end of each session, subjects received six euros for each ECU they won during a subset of the twenty rounds played.[9]

2.2.2 Belief elicitation

The group that did not place bets in treatment 2 was asked before each period to state beliefs about the likelihood that A, as well as B, was the true state in the round. At the time they made their assessments, they had the current odds and the history of previous bets available. Before making his assessments in period i, exactly one of the observers received a private signal with similar content to bettor i's. Each observer received exactly one signal in each round. Each observer stated his beliefs in period i by keying in a vector $\mu_i = (\mu_i^A, \mu_i^B)$, indicating his belief about the probability that the color randomly chosen at the beginning of the round was A or B.[10]

Subjects' assessments were rewarded on the basis of a quadratic scoring rule function. In period i, the payoff when color $j \in \{A, B\}$ was the outcome and μ_i was the belief statement vector was given by € $.15\Big(1 - (\mu_i^k)^2\Big)$ with $k \in \{A, B\}$, $k \neq j$. Thus, if color A was the true state, the greater the weight the subject placed on B, the more subtracted from his endowment of .15 euros. The worst possible statement, placing all weight on the incorrect outcome, yielded a payoff of 0. It can be easily demonstrated that this reward function provides an incentive for risk-neutral subjects to reveal their true beliefs about the probability that each color was chosen.[11]

In the third treatment, each group of eight participants bet against other members of the group in

[8] The belief that the probability of an event decreases when the event has occurred recently, even though the probability of the event is objectively known to be independent across trials, is called the gambler's fallacy. For more discussion of the gambler's fallacy in parimutuel markets, see Terrell (1994).

[9] At the end of the session, the assistant randomly drew the rounds that counted toward participants' earnings. Cubitt, Starmer, and Sugden (1998), in a systematic investigation, find that random lottery designs yield behavior no different from 'single choice' designs in which subjects face just one non-hypothetical task.

[10] In the experiment, subjects entered μ_i^A and μ_i^B as numbers in the interval [0,100], which were described as percentages.

[11] Belief elicitation using a quadratic scoring rule is widely employed in experimental economics (see, among others, Offerman, Sonnemans, and Schram, 1996; Huck and Weizsäcker, 2002; or Nyarko and Schotter, 2002). McKelvey and Page (1990) find that a proper scoring rule successfully induces individuals to reveal their current posterior probabilities of the realization of an exogenous random variable. Although the quadratic scoring rule is not incentive compatible if subjects are not risk neutral or if they fall prey to probability weighting, Sonnemans and Offerman (2001) find that subjects' belief statements are unbiased.

20 repetitions of the betting market. Each participant, in addition to placing bets, stated his beliefs in each period about the likelihood that each color was the true state, in an identical manner to treatment 2. In treatment 3, participants observed only the betting activity of their own group.[12]

3 Models of Individual and Market Behavior

The sequential parimutuel betting market we study in our experiment is a well-defined extensive form game, and in principle, Nash or sequential equilibria can be identified. However, there are two difficulties with using sequential equilibrium or any weaker equilibrium concept as a predictive tool for this market. The first is that taking into account the effect of one's betting decisions on the subsequent actions of multiple future bettors is a demanding calculation unlikely to be consistent with actual participants' behavior. The second is that the large number of equilibria means that the predictions that can be extracted from equilibrium analysis are typically indeterminate.[13]

The *Adaptive Model* described below overcomes these two difficulties. This model assumes that bettors are myopic in the sense that they do not take into account the effect of their own bets on the future decisions of other bettors and the resulting future changes in the odds. That is, each individual acts as if the odds after he makes his choice are the final odds. However, bettors exhibit a high degree of rationality. They maximize a payoff function based on their own information, the actions of previous bettors, and the market odds; and they update their beliefs using Bayes' rule, under the assumption that all previous bettors also fail to take into account the future consequences of their behavior. The Adaptive Model is in the spirit of a growing experimental literature suggesting that expectations are adaptive rather than rational,[14] and generically yields a unique predicted action at each decision node.

3.1 The Adaptive Model

Consider the decision of bettor i, the ith bettor in the sequence. Define the history of bets up to and including bettor i as $s^i = (s_1, \ldots, s_i) \in \{A, B\}^i$. s^i implies that the odds in period i against horse H are

$$O_H(s^i) = \frac{i - h(s^i)}{h(s^i)}.$$

The market odds define a *subjective* (or market) *probability* after period i that horse H wins the race. The subjective probability, P_H^S, is

$$P_H^S(s^i) \equiv \frac{h(s^i)}{i} = \frac{1}{O_H(s^i) + 1}.$$

The subjective probability can be viewed as the implicit price of obtaining a claim to one unit of money in the event that horse H wins the race. In contrast, the *objective probability* after period i that horse H wins the race is

$$P_H^O(q^i) \equiv \Pr(\theta_H \mid q^i),$$

[12] In treatments 1 and 2, three of the twenty rounds were chosen at random at the end of the session to count toward bettors' earnings. In treatment 3, two of the twenty rounds of bets counted toward subjects' earnings. This adjustment served to make the earnings in the three treatments more comparable. All belief submissions counted toward earnings in treatments 2 and 3.

[13] See Koessler and Ziegelmeyer (2003) for a detailed discussion of this point.

[14] For example, Marimon and Sunder (1993) or Johnson, Camerer, Sen, and Rymon (2002) find greater empirical support for adaptive rather than rational expectations. Ziegelmeyer, Broihanne, and Koessler (2004) consider a sequential betting market similar to ours but with symmetrically informed subjects, and find that betting behavior is inconsistent with backward induction. Drehmann, Oechssler, and Roider (2004b) are not able to reject the hypothesis of myopic behavior in information cascade experiments with various forms of payoff externalities.

where $q^i = (q_1, \ldots, q_i)$ is the vector of signals received by all bettors up to and including period i. $P_H^O(q^i)$ is the belief about the probability that horse H wins the race of a hypothetical bettor who would have been able to observe the signals of all bettors up to and including period i. A bettor perceives the payoff of betting on H in period i as

$$u_i(s^i, \theta) = \begin{cases} O_H(s^i) + 1 = \frac{i}{h(s^i)} & \text{if } s_i = H \text{ and } \theta = \theta_H, \\ 0 & \text{if } s_i = H \text{ and } \theta \neq \theta_H. \end{cases}$$

Let k and $\overline{k}$ denote the number of signals favoring horse H and $\overline{H}$, respectively, that bettor i infers from the history of previous bets s^{i-1}. Since pooling can occur at some decision nodes, a particular bet may contain no information about the private signal, and thus $k + \overline{k} < i - 1$ is possible. Accordingly, bettor i's posterior probability that horse H wins the race is

$$\mu_i(\theta_H \mid s^{i-1}; q_i) = \begin{cases} \left(\pi^{k+1} (1-\pi)^{\overline{k}}\right) / \left(\pi^{k+1} (1-\pi)^{\overline{k}} + (1-\pi)^{k+1} \pi^{\overline{k}}\right) & \text{if } q_i = q^H, \\ \left(\pi^{k} (1-\pi)^{\overline{k}+1}\right) / \left(\pi^{k} (1-\pi)^{\overline{k}+1} + (1-\pi)^{k} \pi^{\overline{k}+1}\right) & \text{if } q_i = q^{\overline{H}}. \end{cases} \quad (1)$$

Bettor i bets on horse H if the perceived expected payoff associated with such a bet is greater than the perceived expected payoff associated with betting on horse $\overline{H}$. This is equivalent to saying that bettor i bets on horse H if

$$\mu_i(\theta_H \mid s^{i-1}; q_i) > \frac{h(s^{i-1}) + 1}{i + 1}. \quad (2)$$

Equation (2) implies that the first bettor bets according to his private signal. Given his private signal q_i and after having observed a sequence of previous bets s^{i-1}, each subsequent bettor $i > 1$ bets on either horse A or B according to the following rule. First, bettor i infers the number of signals favoring each horse from s^{i-1}. A signal favoring horse H, q^H, can be inferred from an observed bet s_k in period k if q^H leads to s_k and $q^{\overline{H}}$ leads to $s'_k \neq s_k$. Second, bettor i forms his posterior probability that horse H wins the race using equation (1). Finally, he bets on horse H if inequality (2) holds and on the other horse if it does not hold. In case of indifference, we assume that a bettor bets in accordance with his private information.

The Adaptive Model predicts unique beliefs and strategies for each bettor for all possible histories and profiles of signals. The predictions for all twenty rounds of our three treatments are given in tables 2 and 3, in which objective and subjective probabilities, as well as final odds, are expressed with respect to horse A. The first row in the group of four rows corresponding to each round gives the realization of the signal each bettor $i = 1, \ldots, 8$ receives (the same random draws were used in all sessions). The second row contains the sequence of bets that the Adaptive Model predicts. The third and fourth rows display, respectively, the subjective and objective probabilities. Choices that allow an observer (knowing the bettors' decision rule) or a later mover to infer the private information of the bettor, are represented in bold letters. Choices in which the bettor bets on the favorite and against his private information are classified as *herd behavior* and are indicated with an asterisk. *Contrarian behavior*, betting against own private information and against the favorite, is never predicted.

As an alternative model against which to evaluate the Adaptive Model, we propose a plausible betting rule called the *Private Information Plus Odds* heuristic (*PIPO*). A bettor using this heuristic does not infer any information from past bets but simply responds optimally to the current odds given his private signal. Formally, bettor i bets on horse H if $\Pr(\theta_H \mid q_i) > \frac{h(s^{i-1})+1}{i+1}$. For example, when his private signal is q^A, a bettor complying with the *PIPO* betting rule bets on horse A unless the return from horse B is more than three times the return on horse A.

Table 2: Adaptive Model Predictions for the First Ten Rounds

Round		Period 1	2	3	4	5	6	7	8	Final Odds	Winning Horse
	Signals	q^B	q^B	q^B	q^A	q^A	q^B	q^B	q^B		
	Bets	**B**	**B**	**B**	B^*	B^*	*B*	*B*	*B*		
1	Subjective probability	0	0	0	0	0	0	0	0	Ind	*B*
	Objective probability	.250	.100	.036	.100	.250	.100	.036	.012		
	Signals	q^B	q^A	q^A	q^A	q^A	q^A	q^A	q^A		
	Bets	**B**	**A**	**A**	**A**	*A*	*A*	**A**	*A*		
2	Subjective probability	0	.500	.667	.750	.800	.833	.857	.875	.143	*A*
	Objective probability	.250	.500	.750	.900	.964	.988	.996	.999		
	Signals	q^A	q^B	q^A	q^B	q^A	q^A	q^A	q^A		
	Bets	**A**	**B**	**A**	**B**	**A**	**A**	*A*	*A*		
3	Subjective probability	1	.500	.667	.500	.600	.667	.714	.750	.333	*A*
	Objective probability	.750	.500	.750	.500	.750	.900	.964	.988		
	Signals	q^A	q^A	q^B	q^A	q^A	q^A	q^B	q^A		
	Bets	**A**	**A**	**B**	**A**	*A*	*A*	**B**	**A**		
4	Subjective probability	1	1	.667	.750	.800	.833	.714	.750	.333	*A*
	Objective probability	.750	.900	.750	.900	.964	.988	.964	.988		
	Signals	q^B	q^B	q^A	q^B	q^B	q^A	q^B	q^B		
	Bets	**B**	**B**	**A**	**B**	*B*	B^*	**B**	*B*		
5	Subjective probability	0	0	.333	.250	.200	.167	.143	.125	7	*B*
	Objective probability	.250	.100	.250	.100	.036	.100	.036	.012		
	Signals	q^A	q^A	q^A	q^B	q^B	q^B	q^B	q^B		
	Bets	**A**	**A**	**A**	A^*	A^*	A^*	A^*	A^*		
6	Subjective probability	1	1	1	1	1	1	1	1	0	*B*
	Objective probability	.750	.900	.964	.900	.750	.500	.250	.100		
	Signals	q^B	q^A	q^A	q^B	q^A	q^A	q^A	q^A		
	Bets	**B**	**A**	**A**	**B**	**A**	**A**	*A*	*A*		
7	Subjective probability	0	.500	.667	.500	.600	.667	.714	.750	.333	*A*
	Objective probability	.250	.500	.750	.500	.750	.900	.964	.988		
	Signals	q^B	q^B	q^B	q^B	q^B	q^A	q^A	q^A		
	Bets	**B**	**B**	**B**	*B*	*B*	B^*	B^*	B^*		
8	Subjective probability	0	0	0	0	0	0	0	0	Ind	*B*
	Objective probability	.250	.100	.036	.012	.004	.012	.036	.100		
	Signals	q^A	q^A	q^B	q^A	q^B	q^A	q^A	q^A		
	Bets	**A**	**A**	**B**	**A**	A^*	*A*	**A**	*A*		
9	Subjective probability	1	1	.667	.750	.800	.833	.857	.875	.143	*A*
	Objective probability	.750	.900	.750	.900	.750	.900	.964	.988		
	Signals	q^B	q^B	q^B	q^B	q^B	q^A	q^B	q^B		
	Bets	**B**	**B**	**B**	*B*	*B*	B^*	*B*	*B*		
10	Subjective probability	0	0	0	0	0	0	0	0	Ind	*B*
	Objective probability	.250	.100	.036	.012	.004	.012	.004	.001		

Notes: Boldface—Informative bet; *—Herd behavior.

3.2 Market behavior and measure of information aggregation

To measure the extent to which information aggregation takes place in our markets we use two baselines. The first baseline is a hypothetical parimutuel betting market, in which each bettor has access to all of

Table 3: ADAPTIVE MODEL PREDICTIONS FOR THE LAST TEN ROUNDS

Round		Period 1	2	3	4	5	6	7	8	Final Odds	Winning Horse
	Signals	q^A	q^A	q^B	q^B	q^B	q^A	q^B	q^B		
	Bets	**A**	**A**	**B**	**B**	**B**	**A**	**B**	**B**		
11	Subjective probability	1	1	.667	.500	.400	.500	.429	.375	1.667	*B*
	Objective probability	.750	.900	.750	.500	.250	.500	.250	.100		
	Signals	q^B	q^B	q^A	q^B	q^A	q^A	q^A	q^A		
	Bets	**B**	**B**	**A**	**B**	B^*	B^*	**A**	**A**		
12	Subjective probability	0	0	.333	.250	.200	.167	.286	.375	1.667	*A*
	Objective probability	.250	.100	.250	.100	.250	.500	.750	.900		
	Signals	q^B	q^A	q^B	q^B	q^B	q^B	q^B	q^B		
	Bets	**B**	**A**	**B**	**B**	*B*	*B*	**B**	*B*		
13	Subjective probability	0	.500	.333	.250	.200	.167	.143	.125	7	*B*
	Objective probability	.250	.500	.250	.100	.036	.012	.004	.001		
	Signals	q^A	q^A	q^A	q^A	q^B	q^A	q^B	q^A		
	Bets	**A**	**A**	**A**	*A*	A^*	*A*	A^*	*A*		
14	Subjective probability	1	1	1	1	1	1	1	1	0	*A*
	Objective probability	.750	.900	.964	.988	.964	.988	.964	.988		
	Signals	q^B	q^B	q^B	q^A	q^B	q^B	q^B	q^B		
	Bets	**B**	**B**	**B**	B^*	*B*	*B*	*B*	*B*		
15	Subjective probability	0	0	0	0	0	0	0	0	Ind	*B*
	Objective probability	.250	.100	.036	.100	.036	.012	.004	.001		
	Signals	q^B	q^B	q^B	q^B	q^B	q^B	q^A	q^A		
	Bets	**B**	**B**	**B**	*B*	*B*	*B*	B^*	B^*		
16	Subjective probability	0	0	0	0	0	0	0	0	Ind	*B*
	Objective probability	.250	.100	.036	.012	.004	.001	.004	.012		
	Signals	q^A	q^A	q^A	q^A	q^B	q^B	q^A	q^A		
	Bets	**A**	**A**	**A**	*A*	A^*	A^*	*A*	*A*		
17	Subjective probability	1	1	1	1	1	1	1	1	0	*A*
	Objective probability	.750	.900	.964	.988	.964	.900	.964	.988		
	Signals	q^A	q^B	q^B	q^A	q^B	q^B	q^B	q^B		
	Bets	**A**	**B**	**B**	**A**	**B**	**B**	*B*	*B*		
18	Subjective probability	1	.500	.333	.500	.400	.333	.286	.250	3	*B*
	Objective probability	.750	.500	.250	.500	.250	.100	.036	.012		
	Signals	q^B	q^A	q^A	q^B	q^B	q^B	q^B	q^B		
	Bets	**B**	**A**	**A**	**B**	**B**	**B**	*B*	*B*		
19	Subjective probability	0	.500	.667	.500	.400	.333	.286	.250	3	*B*
	Objective probability	.250	.500	.750	.500	.250	.100	.036	.012		
	Signals	q^A	q^A	q^A	q^B	q^A	q^A	q^A	q^B		
	Bets	**A**	**A**	**A**	A^*	*A*	*A*	*A*	A^*		
20	Subjective probability	1	1	1	1	1	1	1	1	0	*A*
	Objective probability	.750	.900	.964	.900	.964	.988	.996	.988		

Notes: Boldface—Informative bet; *—Herd behavior.

the private information of previous bettors, including his own. This baseline represents the maximum possible level of information aggregation the market could achieve. The dynamics of prices in such a market are straightforward to characterize: bettor i bets on horse H if $\Pr(\theta_H \mid q^i) > \frac{h(s^{i-1})+1}{i+1}$. The resulting final *observable signals* (OS) subjective probability of H given the signal profile q^n is denoted by $P_H^{OS}(q^n)$.

As an alternative baseline, consider a hypothetical parimutuel betting market where bettors simply bet randomly. No information aggregation takes place in the market as bettors completely ignore both their private information and the pricing mechanism. In such a market the subjective probability of each outcome equals one half for any signal profile.

Now define $V = \sum_{q^8 \in Q_E^8} \frac{1}{20} \left| P_A^{OS}(q^8) - P_A(q^8) \right|$, and $V_{\max} = \sum_{q^8 \in Q_E^8} \frac{1}{20} \left| P_A^{OS}(q^8) - 1/2 \right|$, where q^8 denotes a signal profile ($n = 8$), $P_A(q^8)$ denotes the observed subjective probability for horse A after period 8, and Q_E^8 denotes the set of twenty private information profiles used in the experiment. $V_{\max}$, the difference between the maximum and the minimum possible level of information aggregation, is a measure of the maximum possible amount of information that can be aggregated and V is interpreted as a measure of the amount of information that the market has not aggregated. We use the average observed value of $IA = (V_{\max} - V) / V_{\max}$ as a normalized measure of how much information aggregation occurs in our markets. IA measures the amount of information actually aggregated as a fraction of the maximum possible.

4 Results

Figure 1 plots the time path of horse A's objective probability and subjective probability for each round of the experiment. A separate time series is provided for each of the three treatments in each round, and the data are averaged over the sessions within each treatment. There are several general patterns apparent at first glance. In all treatments the correct horse is usually the favorite at the end of the betting process, as the observed subjective probability typically implies the same favorite as the objective probability. The only exceptions are in rounds 6 and 12 for treatment 1, and rounds 6, 11, and 19 for treatments 2 and 3. The final subjective probability in treatment 3 is usually (in 15 of 20 rounds) closer to the objective probability than in each of the other treatments, suggesting better information aggregation in treatment 3. There does not seem to be a systematic difference with respect to information aggregation between treatments 1 and 2, as average final subjective probabilities in treatment 2 are closer to objective probabilities than in treatment 1 for 12 of 20 rounds. Deviations of the subjective from the objective probability are usually in the direction of one-half in all treatments, which is consistent with the presence of a favorite-longshot bias.

Figure 2 plots the average[15] time path of horse A's objective probability compared to the average belief statements in both treatments 2 and 3. It is evident from the figure that the differences between the two treatments are considerable. Belief statements in treatment 3 are on average closer to the objective probability than in treatment 2 in all 20 rounds. Average beliefs in treatment 2 are closer to 1/2, an equal probability on each outcome, than in treatment 3 in 17 of 20 rounds. The exceptions are rounds 6, 11, and 19, the only ones in which the subjective probabilities on betting behavior yields an incorrect prediction in a majority of treatments.

In the following subsections we study the data in more detail. In subsection 4.1, we compare the information aggregation properties of the three treatments to each other and to theoretical benchmarks. In subsection 4.2, we compare betting behavior to the predictions of the Adaptive Model and the $PIPO$ heuristic, and in subsection 4.3, we study belief statements as well as the consistency between belief statements and the subjective and objective probabilities.

[15]Because in each period only one subject among the eight who state their beliefs is endowed with a private signal, in order to make fair comparisons between the belief statements and the objective probability of horse A, we compute the *average* objective probability in period i as $\left(7\Pr(\theta_A \mid q^{i-1}) + \Pr(\theta_A \mid q^i)\right)/8$.

Figure 1: Horse A's Objective Probability versus Horse A's Observed Subjective Probability (average across all groups)

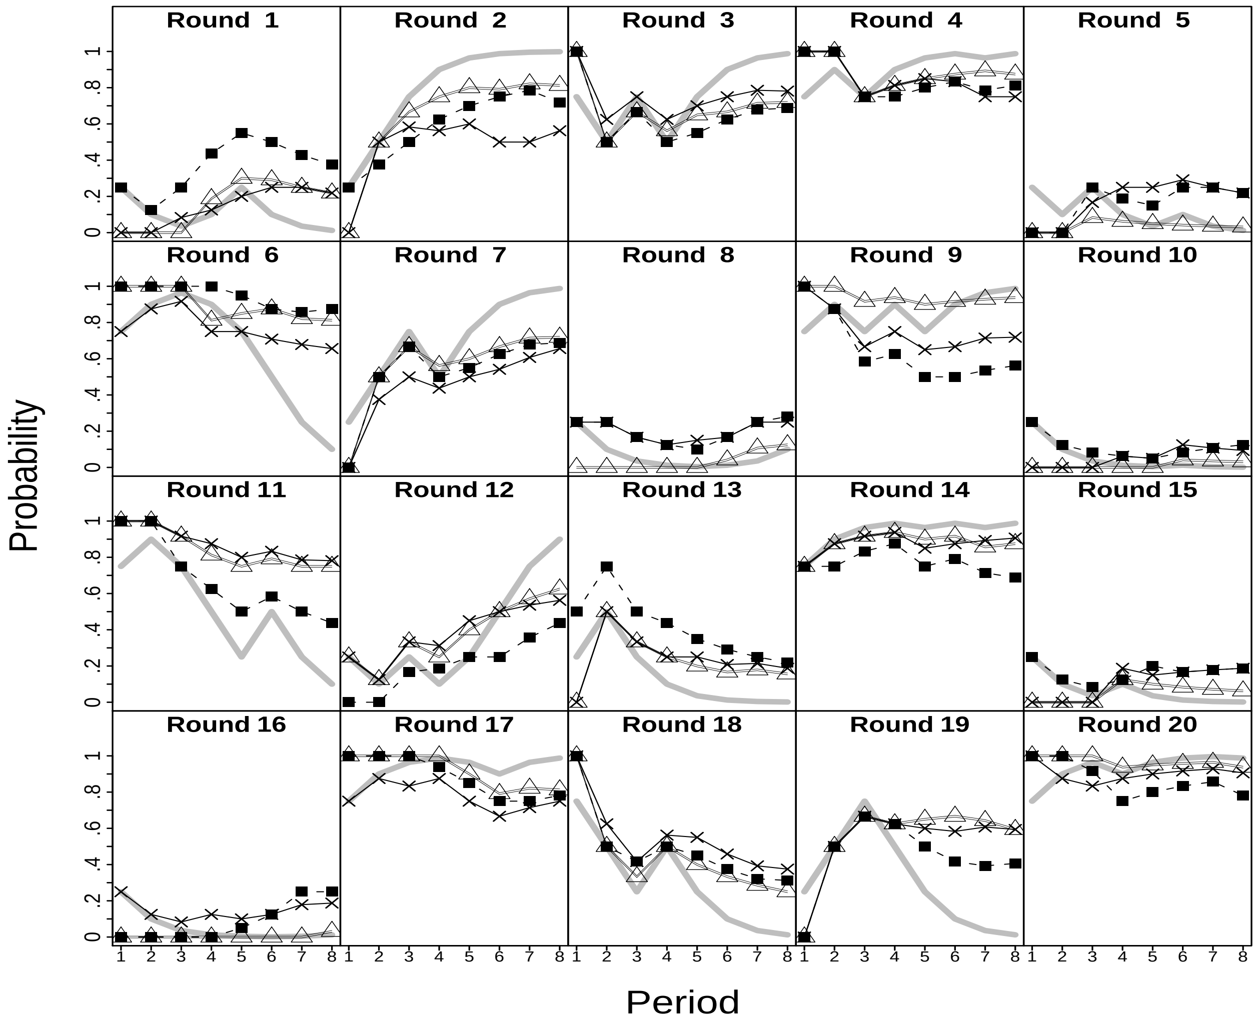

Notes: Objective probability; Subjective probability in treatment 1; Subjective probability in treatment 2; Subjective probability in treatment 3.

4.1 Market behavior: Prices and information aggregation

We begin our analysis at the market level. We consider the ability of the markets to aggregate information by treatment and in relation to the Adaptive Model. Our results are listed below, and statistical support is provided in the paragraphs that follow.

RESULT 1: INFORMATION AGGREGATION. *Elicitation of beliefs of bettors improves the information aggregation ability of the market. Deviations are consistent with a favorite-longshot bias in the first two treatments. The bias is sharply reduced when bettors' beliefs are elicited. The level of information aggregation is lower than under the Adaptive Model.*

SUPPORT: The average value of IA for each of the four groups of bettors that participated in treatment 3 is higher than the average value of IA of any group in treatments 1 and 2. Thus, a conservative rank-sum test, using each group as an observation, rejects the hypotheses at the 5 percent level that (a) treatments 1 and 3, and (b) treatments 2 and 3, aggregate the same amount of information. A similar test for a difference between treatments 1 and 2 is insignificant at conventional levels.

Figure 2: Average Belief Statements versus Horse A's Objective Probability

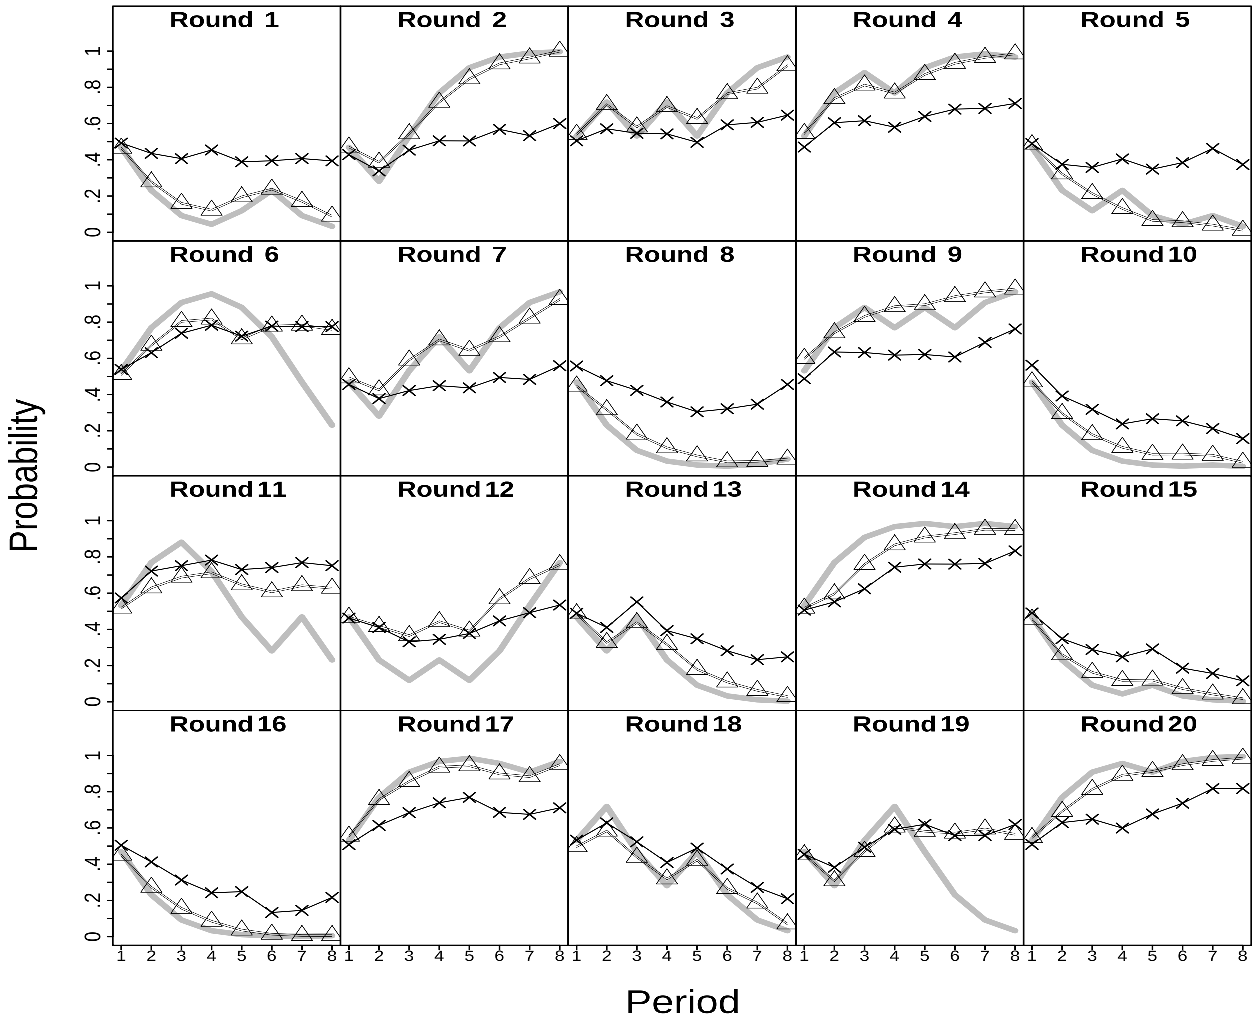

Notes: Objective probability;
Average beliefs in treatment 2; Average beliefs in treatment 3.

By the end of every round, the objective probability for one of the horses equals at least 0.9. The final subjective probabilities are always divisible by 1/8, so that subjective probabilities equal to 1 or 7/8 in favor of the objective favorite do not indicate a deviation from the objective probability, but any subjective probability of .75 or lower on the favorite is consistent with a favorite-longshot bias. This occurs in 87.50% of rounds in treatment 1, 68.75% in treatment 2, but only 43.75% in treatment 3.[16]

However, we also observe that information aggregation, even in treatment 3, is below the level under the Adaptive Model. The values of IA observed in treatments 1 and 2, 50.45% and 53.18%, are approximately equal to the 50.91% under the $PIPO$ heuristic. However, they are lower than the 68.18% in treatment 3, which is in turn considerably (and significantly) lower than the 87.27% under the Adaptive Model (p-value $< .001$ according to a mixed effects model where IA is a linear function of the fixed effect

[16] The percentages are 73.85%, 60.32%, and 33.82% respectively for treatment 1, 2, and 3 if the few rounds where the observed subjective probability does not imply the same favorite as the objective probability are excluded. Of course, we recognize that in our setting with only two alternatives, any source of errors, including pure randomness in decision making, will tend to produce a favorite-longshot bias. We provide a detailed analysis of the errors in decision making in our data in subsection 4.2.

treatment and the 12 independent groups are modeled as random effects).[17] □

4.2 Individual behavior: Bets placed

In this section, we consider the extent to which individual bets are consistent with the Adaptive Model and the $PIPO$ heuristic, and how the level of consistency varies by treatment. Our results are summarized in the following statements.

RESULT 2: BETTING BEHAVIOR and the ADAPTIVE MODEL. *Individual bets in all treatments provide strong support for the Adaptive Model. The formulation of both bets and belief statements improves the predictive performance of the Adaptive Model. In all treatments, the Adaptive Model correctly predicts more individual betting decisions than the $PIPO$ heuristic.*

SUPPORT: We compare actual betting decisions with the predictions of the Adaptive Model and of the $PIPO$ heuristic.[18] The predictions of these two benchmarks coincide in most instances, but in 23.91% of the observations in treatment 1, 26.09% in treatment 2, and 33.59% in treatment 3, the Adaptive Model and the $PIPO$ heuristic make different predictions. The second row of table 4 shows the percentage of bets that are consistent with the predictions of the Adaptive Model. A large majority of the actual bets is compatible with the Adaptive Model in each treatment. A greater percentage is consistent with the Adaptive Model in treatment 3 than in treatments 1 and 2, between which no difference is observed. The Adaptive Model is therefore more relevant for predicting decisions made when bettors explicitly state their beliefs.

The fourth row of data in table 4 gives the percentage of bets consistent *only* with the predictions of the Adaptive Model in those instances in which it and the $PIPO$ heuristic make different predictions. In all three treatments a significant majority of the participants' decisions are compatible with the Adaptive Model predictions. A binomial test rejects the hypothesis that the percentage compatible with the Adaptive Model is less than or equal to 1/2 at the one percent level in each of the three treatments. The Adaptive Model thus organizes our data better than the $PIPO$ heuristic in all three treatments, and the difference is especially pronounced in treatment 3.

The last row of the table illustrates the combined predictive power of the two benchmarks. It shows the relative frequency of subjects' bets which are inconsistent with both the Adaptive Model and with the $PIPO$ heuristic. The fact that, in each treatment, approximately half of the bets that are inconsistent with the Adaptive Model are also inconsistent with the $PIPO$ heuristic indicates that $PIPO$ is of no additional value as a predictor over and above the Adaptive Model (since predicting bets in accordance or alternatively counter to the $PIPO$ heuristic for observations counter to the Adaptive model yields essentially the same hit rate). □

Result 2 establishes that requiring participants to state their beliefs leads to bets that are closer to the Adaptive Model predictions. One possible explanation for this pattern is that belief elicitation reduces noise in decision-making, while another possibility is that it systematically reduces particular types of deviations from the Adaptive Model, perhaps because the placement of bets without belief elicitation may cause an overemphasis on the odds against outcomes rather than on the probability of their occurrence. Result 3 suggests that the latter is the case.

[17] The average absolute differences between the subjective and objective probabilities are .180, .174, and .138 in treatments 1, 2, and 3, respectively. For the data from the last period only, the average absolute differences are .299, .280, and .210 in treatments 1 - 3. For comparison, the average differences predicted under the Adaptive Model over all periods equals .129, and for the last period equals .176.

[18] If the previous sequence of decisions is not possible under the Adaptive Model for any profile of signals, we assume that it is common knowledge that any bettor who deviates follows his private signal. In all of our analysis, we obtain very similar results if we assume instead that bettors' deviations are uncorrelated with their private signals.

Table 4: Percentage of Decisions Consistent with the Adaptive Model and/or the *PIPO* heuristic

	Treatment 1	Treatment 2	Treatment 3
Percentage of bets compatible with the Adaptive Model	81.09% (519/640)	81.25% (520/640)	87.03% (557/640)
Percentage of bets compatible with the *PIPO* heuristic	75.62% (487/640)	68.28% (436/640)	65.94% (422/640)
Percentage of bets compatible with the Adaptive Model in instances in which the Adaptive Model and the *PIPO* heuristic make different predictions	61.44% (94/153)	74.85% (125/167)	81.40% (175/215)
Percentage of bets neither compatible with the Adaptive Model nor with the *PIPO* heuristic	9.69% (62/640)	12.19% (78/640)	6.72% (43/640)

RESULT 3: EFFECT of BELIEF ELICITATION on DEVIATIONS from the ADAPTIVE MODEL. *Elicitation of beliefs almost completely eliminates contrarian betting and reduces the incidence of the failure to herd. However, it does not reduce the incidence of incorrect herding.*

SUPPORT: Deviations from the Adaptive Model can be classified as one of three types: i) Contrarian behavior: betting on the longshot and against one's own private information[19]; ii) Incorrect herding: incorrectly betting on the favorite and against one's own private information; and iii) Failure to herd: incorrectly betting according to one's own private information and against the favorite. The left side of table 5 reports, for each treatment, the percentage of instances in which each type of deviation occurred when it was possible. Only contrarian behavior and failure to herd are less frequent in treatment 3 than in the other two treatments. Contrarian behavior decreases by 71% from the treatment 1 level and by 73% from the treatment 2 level. On the other hand, the incidence of incorrect herding is somewhat higher in treatment 3 than in the other two treatments. □

Table 5: Percentage of Instances of the Occurrence of Different Types of Deviations from the Adaptive Model (in the Data and under *QRAM* for Different Values of Lambda)

	Treatment			*QRAM*'s precision parameter				
	1	2	3	10	7	5	3	.1
Contrarian behavior	10.55% (436)	11.49% (444)	3.10% (452)	0.28% (447)	1.12% (444)	3.66% (439)	11.19% (432)	49.4% (388)
Incorrect herding	22.13% (122)	29.52% (105)	34.38% (87)	37.60% (82)	34.82% (87)	33.20% (97)	31.94% (115)	56.06% (157)
Failure to herd	58.54% (82)	41.76% (91)	38.61% (101)	32.85% (111)	35.66% (109)	40.08% (105)	50.65% (93)	50.60% (95)

Note: Numbers in parentheses are the frequencies of *potential* deviations.

Contrarian behavior is on average the most costly of the three types of deviation from the Adaptive Model for the bettor. On average, other types of deviations are 25.71%, 26.32%, and 33.33% less costly in treatment 1, 2, and 3, respectively. In general, belief elicitation reduces costly deviations from the Adaptive Model by a greater percentage than less costly deviations. For an expected cost greater than or

[19]We include in this class betting against one's private information when the current odds on each outcome are equal.

equal to .8 currency units, the relative frequency of deviations is 12.59%, 13.99%, and 4.10% in treatments 1, 2, and 3, respectively, while for an expected cost less than .3, the relative frequency of deviations is 37.78%, 41.03%, and 44.19% in treatments 1, 2, and 3, respectively.

The fact that the deviations from the Adaptive Model are less likely the costlier they are, as the data in the last paragraph indicates, suggests that a statistical extension of the Adaptive Model, incorporating the assumption that errors are less likely to occur the more costly they are, might provide a good fit to the data. According to the Adaptive Model, bettor i's (perceived) expected utility from betting on horse H is $U_i(H \mid s^{i-1}; q_i) = \mu_i(\theta_H \mid s^{i-1}; q_i)\frac{i}{1+h(s^{i-1})}$, where $\mu_i(\theta_H \mid s^{i-1}; q_i)$ is the bettor's belief that horse H will win the race given the information available. One specification for incorporating errors into the best response of bettors is to assume that bettor i bets on horse H with probability

$$\Pr(i \text{ bets on } H) = \frac{e^{\lambda U_i(H|s^{i-1};q_i)}}{e^{\lambda U_i(A|s^{i-1};q_i)} + e^{\lambda U_i(B|s^{i-1};q_i)}}, \tag{3}$$

where errors are measured in terms of a precision parameter $\lambda \geq 0$. If $\lambda = 0$, then bettors are betting entirely randomly, while for $\lambda = \infty$, bets perfectly match the Adaptive Model's predictions.

We consider several statistical generalizations of the Adaptive Model that incorporate the logistic decision rule (3).

1. $QRAM$: We assume that each bettor knows the decision rules of previous bettors. We refer to this statistical model as the Quantal Response Adaptive Model ($QRAM$) to emphasize the close similarity to Quantal Response Equilibrium (see McKelvey and Palfrey, 1998).

2. Step-k-$QRAM$ ($k = 1, \ldots, 7$): This is a further generalization of $QRAM$ in which previous bettors' strategies are correctly deduced only up to a given depth of reasoning (see Camerer, Ho, and Chong, 2004, or Kübler and Weizsäcker, 2004). Step-1-$QRAM$ is a model in which each bettor uses the logistic decision rule (3) but believes that all previous bettors bet randomly (or, equivalently, that they use the logistic decision rule with $\lambda = 0$). More generally, in the step-$k+1$-$QRAM$ model each bettor uses the logistic decision rule (3) and believes that all previous bettors bet according to the step-k-$QRAM$.[20] Notice that these models include several interesting heuristics as special cases. When the precision parameter λ tends to infinity, the step-1-$QRAM$ is equivalent to the $PIPO$ heuristic. When the depth of reasoning equals two and λ tends to infinity, decisions are a best response to the $PIPO$ heuristic.

3. NAM: This third model of behavior assumes that bettors use the logistic decision rule (3), and believe that other bettors use the Adaptive Model while ignoring the fact that they also make errors (or, equivalently, that they use the logistic decision rule with $\lambda = \infty$). This model, which we call the Noisy Adaptive Model (NAM), also spans a range of behavior for which the Adaptive Model and random decisions are special cases.

Our findings concerning the predictive power of the different statistical models considered are summarized in the statement of the next result.

RESULT 4: COMPARISON of STATISTICAL GENERALIZATIONS of the ADAPTIVE MODEL. *Contrary to NAM, QRAM captures the pattern of deviations from the Adaptive Model observed across the three treatments. QRAM outperforms step-k-QRAM in predictive power, so that assuming limited depth of reasoning does not improve predictions. The belief elicitation procedure reduces the estimated incidence of commonly known errors as measured by the parameter* λ.

SUPPORT: Table 6 reports the results of the maximum-likelihood estimation of $QRAM$, for the three treatments. The table also contains (in parentheses) the standard error of each estimated parameter and the marginal level of significance for the parameter. The negative log likelihood is given in the last row.

[20] Because in our setting at most eight steps of reasoning can be considered, step-8-$QRAM$ is equivalent to $QRAM$.

Table 6: *QRAM*'s Precision Parameter Estimated from Experimental Data

	Treatment 1	Treatment 2	Treatment 3
λ	2.890	2.837	7.184
	(.193, .000)	(.257, .000)	(.590, .000)
$-ll^*$	277.2882	301.7246	187.9936

Notes: Numbers in parentheses are (i) the standard error of the parameter, and (ii) the marginal level of significance that the parameter differs from zero. $-ll^*$ is the negative log likelihood.

Likelihood ratio tests indicate that the precision parameter λ is not significantly different between treatments 1 and 2, but it is significantly greater in treatment 3 than in the other two treatments. This lower level of error is consistent with better predictive power of the Adaptive Model in treatment 3 than in the other two treatments. The log likelihood is much larger in treatment three than in the two first treatments.

To verify whether the difference in the precision parameter between the first two treatments and the third one can account for the pattern of deviations from the Adaptive Model reported on the left side of table 5, we generate sequences of bets according to *QRAM* for numerous values of λ and then we compute the relative frequencies of each type of deviation.[21] The results, reported on the right side of table 5, show that contrarian behavior and failure to herd decline as λ increases (contrarian behavior declines very rapidly), and that the model allows for the possibility that incorrect herding deviations *increase* for a *smaller* error rate. Clearly, *NAM* cannot account for the fact that some deviations from the Adaptive Model increase when the level of errors decreases, and is therefore inadequate in explaining the impact of the belief elicitation procedure on bettors' behavior.

QRAM outperforms step-*k*-*QRAM* since it generates a higher maximized likelihood in each treatment. Increasing the depth of reasoning from one to four levels improves the predictive power considerably in all treatments. In fact, step-4-*QRAM* and *QRAM* maximized likelihoods are almost identical, as for precision parameter values larger than one both models lead to almost identical predictions. While we cannot conclude that the "rational expectations" assumption is supported, subjects' depth of reasoning is considerable in our markets. □

4.3 Individual behavior: Belief statements

In this section, we consider the consistency between bets and belief statements in treatments 2 and 3. We compare the belief statements to those under *QRAM*. We also evaluate whether placing bets helps participants to correctly guess the winning horse by comparing the discrepancy between belief statements and objective probabilities in treatments 2 and 3. Our results are listed below.

RESULT 5: CONSISTENCY between BETTING DATA and BELIEF STATEMENTS. *When the same individuals both place bets and state beliefs, there is greater consistency between bets and belief statements than when different individuals state beliefs and place bets.*

SUPPORT: We measure the degree of consistency between the betting data and the belief statements using the average absolute difference between individuals' belief statements and the subjective probabili-

[21] For each of the twenty rounds, one hundred sequences of bets were generated for each precision parameter. Thus, the amount of simulated data, on which the analysis of the deviations from the Adaptive Model is based, is 25 times larger than the amount of experimental data for a given treatment.

ties. All four groups in treatment 3 have lower absolute differences than any group in treatment 2. Thus, a rank-sum test, using each group as an observation, rejects the hypothesis at the 5 percent level that the degree of consistency between belief statements and market estimates in treatments 2 and 3 are equal.

Under the assumption that bets and beliefs are generated by the Adaptive Model, we also obtain the result that bets and belief statements display a greater degree of consistency in treatment 3 than in treatment 2. Table 7 shows the average value of the absolute difference between the belief statements and the beliefs predicted by $QRAM$ for each group stating beliefs in treatments 2 and 3. The predicted beliefs are averages from simulations generated from the estimates given in table 6. The overall difference between the stated and the simulated beliefs is almost twice as large in treatment 2 as in treatment 3. The average difference for each of the four groups of bettors that participated in treatment 3 is lower than any of the average differences in treatment 2.

Table 7: Average Absolute Difference between Belief Statements and $QRAM$'s Beliefs

	Treatment 2	Treatment 3
Group 1	.1815	.0868
Group 2	.1771	.0841
Group 3	.1583	.0868
Group 4	.1549	.1100
Overall	.1679	.0919

□

A comparison of $QRAM$, step-k-$QRAM$, and NAM indicates that $QRAM$ leads to the lowest value of the absolute difference between simulated beliefs and the belief statements in treatment 3. For treatment 2, however, the lowest value is achieved by the step-3-$QRAM$ for a precision parameter equal to 1.13. Indeed, the distance corresponds to an average absolute difference of .1454 and is lower under step-3-$QRAM$ than under $QRAM$ for all groups of bettors in treatment 2. This discrepancy between belief statements and bets placed in treatment 2 suggests that observers believe that the bettors use fewer steps of reasoning than they actually do.

RESULT 6: EFFECT of PLACING BETS on ACCURACY of BELIEF STATEMENTS. *Placing bets improves the accuracy of belief statements.*

SUPPORT: Table 8 shows the average value of the absolute difference between the average belief statements and the objective probability for each group participating in treatments 2 and 3. The average difference for each of the four groups of bettors that participated in treatment 3 is lower than any of the average differences in treatment 2, whether all periods or only the final periods are used in the calculation. In addition, subjects' belief statements in treatment 3 lead to payoffs that are 14.68% higher than those in treatment 2. If the comparison between treatments is restricted to instances in which the previous sequences of bets and private signals are identical, the difference is 10.46%.[22]

□

5 Conclusion

Both financial markets and parimutuel betting markets are characterized by numerous investors with heterogeneous information who seek to profit through trading as their information becomes aggregated

[22] Although the difference in payoffs is small, it reflects a substantial difference in belief statements. A well-known property of the quadratic scoring rule is that the payoff function is quite flat in the neighborhood of the optimum.

Table 8: Average Absolute Difference between Belief Statements and Objective Probability

	Treatment 2		Treatment 3	
	all periods	final period	all periods	final period
Group 1	.2025	.2899	.0901	.0915
Group 2	.2225	.3354	.0652	.0833
Group 3	.1963	.2975	.0951	.1292
Group 4	.1875	.2874	.1012	.1292
Overall	.2022	.3025	.0878	.0908

in the prices. The formation of individual beliefs and trading behavior in asset markets is a complex phenomenon whose main features remain to be identified. This study investigates betting behavior and market outcomes in experimental parimutuel betting markets by focusing on the influence of a) the current market price, b) the investors' private information, c) the investors' beliefs about others' degree of rationality, and d) the history of prior bets. Three treatments are considered: one in which individuals only place bets, a second in which betting data are made available to observers who state beliefs about the likelihood of each outcome, and a third in which individuals both place bets and state their beliefs about the likelihood of each outcome. We find that belief elicitation significantly alters betting behavior and improves the ability of the market to aggregate private information in pricing patterns. We also observe that belief statements are both more accurate and more consistent with market activity in the third treatment. However, betting decisions are no better when they are observed than when they are not observed, if beliefs of bettors themselves are not elicited.

The Adaptive Model describes well the overall patterns we observe. However, a statistical generalization of the model provides greater accuracy in prediction and provides a plausible framework for understanding the influence of the belief elicitation procedure on betting behavior. On the one hand, belief elicitation improves decision making by reducing the failure to herd and contrarian behavior, both of which involve betting on the longshot when it is not optimal to do so. The role of contrarian behavior in distorting prices away from full information levels is also emphasized by Cipriani and Guarino (2005), and Drehmann, Oechssler, and Roider (2004a), who study informational cascades in laboratory financial markets. A simple heuristic of trading in accordance with one's own private information is always optimal in their setting, but nonetheless they observe a high incidence of contrarian behavior. In our experimental betting markets, contrarian behavior generates a favorite-longshot bias, which is sharply reduced when bettors are required to submit beliefs. On the other hand, belief elicitation does not reduce the incidence of incorrect herding which is somewhat higher in treatment 3 than in the other two treatments. Still, our results on herding behavior are rather in line with the predictions of the Adaptive Model in all treatments as failure to herd and incorrect herding compensate each other. If anything, too little herding behavior is observed. We therefore confirm another finding of Cipriani and Guarino (2005), and Drehmann, Oechssler, and Roider (2004a), that irrational herding does not seem to be an important force in the presence of a flexible market price.

Experimental economists have found that eliciting beliefs of subjects is a useful methodological tool for theory testing and for gaining insight into human decision making. Indeed, Manski (2004) argues that applied economic research more generally can benefit from combining choice data with self-reports of expectations, elicited in the form of subjective probabilities, to understand economic behavior. Our results suggest yet another use for belief elicitation procedures: as potential instruments to improve a market's capacity to aggregate information.

References

Anderson, L. R., and C. A. Holt (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87(5), 847–862.

Boun My, K. (2002): "A VB Application for Sequential Betting with Asymmetric Information," CNRS, BETA-Theme.

Camerer, C., T.-H. Ho, and J.-K. Chong (2004): "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119(3), 861–898.

Cipriani, M., and A. Guarino (2005): "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, forthcoming.

Costa-Gomes, M. A., and G. Weizsäcker (2005): "Stated Beliefs and Play in Normal-Form Games," mimeo.

Croson, R. T. A. (2000): "Thinking Like a Game Theorist: Factors Affecting the Frequency of Equilibrium Play," *Journal of Economic Behavior and Organization*, 41(3), 299–314.

Cubitt, R. P., C. Starmer, and R. Sugden (1998): "On the Validity of the Random Lottery Incentive System," *Experimental Economics*, 1(2), 115–131.

Drehmann, M., J. Oechssler, and A. Roider (2004a): "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment," *American Economic Review*, forthcoming.

——— (2004b): "Herding with and without Payoff Externalities - An Internet Experiment," Bonn Econ Discussion Paper No. 15/2004, University of Bonn.

Griffith, R. M. (1949): "Odds Adjustments by American Horse-Race Bettors," *American Journal of Psychology*, 62, 290–294.

Huck, S., and G. Weizsäcker (2002): "Do players correctly estimate what others do: Evidence of conservatism in beliefs," *Journal of Economic Behavior and Organization*, 47(1), 71–85.

Hung, A., and C. Plott (2001): "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions," *American Economic Review*, 91(5), 1508–1520.

Johnson, E. J., C. Camerer, S. Sen, and T. Rymon (2002): "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *Journal of Economic Theory*, 104(1), 16–47.

Jullien, B., and B. Salanié (2000): "Estimating Preferences under Risk: The Case of Racetrack Bettors," *Journal of Political Economy*, 108(3), 503–530.

Koessler, F., and A. Ziegelmeyer (2003): "Parimutuel Betting under Asymmetric Information," Discussion Paper on Strategic Interaction No. 34-2003, Max Planck Institute for Research into Economic Systems, Jena.

Kübler, D., and G. Weizsäcker (2004): "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *Review of Economic Studies*, 71(2), 425–441.

Manski, C. F. (2004): "Measuring Expectations," *Econometrica*, 72(5), 1329–1376.

Marimon, R., and S. Sunder (1993): "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," *Econometrica*, 61(5), 1073–1107.

McKelvey, R. D., and T. Page (1990): "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58(6), 1321–1339.

McKelvey, R. D., and T. R. Palfrey (1998): "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, 1(1), 9–41.

Nyarko, Y., and A. Schotter (2002): "An Experimental Study of Belief Learning Using Elicited Beliefs," *Econometrica*, 70(3), 971–1005.

Offerman, T., J. Sonnemans, and A. Schram (1996): "Value Orientations, Expectations and Voluntary Contributions in Public Goods," *Economic Journal*, 106(437), 817–845.

Ottaviani, M., and P. N. Sorensen (2004): "The Timing of Bets and the Favorite-Longshot Bias," mimeo.

——— (2005): "Noise, Information and the Favorite-Longshot Bias," mimeo.

Plott, C. R., J. Witt, and W. C. Yang (2003): "Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results," *Economic Theory*, 22(2), 311–351.

Rutström, E. E., and N. Wilcox (2004): "Learning and Belief Elicitation: Observer Effects," mimeo.

Sauer, R. D. (1998): "The Economics of Wagering Markets," *Journal of Economic Literature*, 36(4), 2021–2064.

Schnytzer, A., and Y. Shilony (1995): "Inside Information in a Betting Market," *The Economic Journal*, 105(431), 963–971.

Sonnemans, J., and T. Offerman (2001): "Is the Quadratic Scoring Rule really incentive compatible?," Working paper, CREED, University of Amsterdam.

Terrell, D. (1994): "A Test of the Gambler's Fallacy: Evidence from Pari-mutuel Games," *Journal of Risk and Uncertainty*, 8, 309–317.

Thaler, R., and W. T. Ziemba (1988): "Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries," *Journal of Economic Perspectives*, 2(2), 161–174.

Vaughan Williams, L. (1999): "Information Efficiency in Betting Markets: A Survey," *Bulletin of Economic Research*, 51(1), 1–30.

Warglien, M., M. G. Devetag, and P. Legrenzi (1998): "Mental Models and Naive Play in Normal Form Games," Working paper, Universita Ca' Foscari di Venezia.

Wilcox, N., and N. Feltovich (2001): "Thinking Like a Game Theorist: Comment," mimeo.

Ziegelmeyer, A., M.-H. Broihanne, and F. Koessler (2004): "Sequential Parimutuel Betting in the Laboratory," *Journal of Risk and Uncertainty*, 28(2), 165–186.