

Information Aggregation and Beliefs in Experimental Parimutuel Betting Markets*

Frédéric KOESSLER, Charles NOUSSAIR, and Anthony ZIEGELMEYER[†]

June 2007

Abstract

We study sequential parimutuel betting markets with asymmetrically informed bettors, using an experimental approach. In one treatment, groups of eight participants play twenty repetitions of a sequential betting game. The second treatment is identical, except that bettors are observed by other participants who assess the winning probabilities of each potential outcome. In the third treatment, the same individuals make bets and assess the winning probabilities of the outcomes. A favorite-longshot bias is observed in the first and second treatments, but does not exist in the third treatment. Information aggregation is better in the third than in the other two treatments, and contrarian betting is almost completely eliminated by the belief elicitation procedure. Making bets improves the accuracy of stated beliefs. We propose a theoretical model, the *Adaptive Model*, to describe individual behavior and we find that it effectively explains betting decisions, especially in the third treatment.

KEYWORDS: Parimutuel betting; Information aggregation; Elicited beliefs; Experimental economics.

JEL CLASSIFICATION: C72; C92; D82.

*We thank Dan Levin and participants at seminars in Bonn, Cergy-Pontoise, Jena, Paris, at the GEW-Tagung in Wittenberg, at the AFFI International Conference 2004, at the 2005 London Mini-Conference on Information and Prediction Markets, 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, Faculty of Economics and Business Administration, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands. Email: C.N.Noussair@uvt.nl. Ziegelmeyer: Max Planck Institute of Economics, Strategic Interaction Group, Kahlaische Strasse 10, D-07745 Jena, Germany. Email: ziegelmeyer@econ.mpg.de.

1 Introduction

Parimutuel betting markets are of considerable empirical importance. Betting turnover on horse racing in the United States alone in 2005 was approximately 14 billion US dollars (International Federation of Horseracing Authorities, <http://www.horseracingintfed.com>). However, parimutuel betting markets are mainly of interest in the economics and finance research communities because they can be viewed as simple representations of financial markets (Thaler and Ziemba, 1988; Sauer, 1998; Vaughan Williams, 1999). In parimutuel win-betting markets, individuals bet money on an outcome, such as a horse winning a race. If the horse wins they get to keep their bet as well as share with the other winners all the money placed on losing horses (less market maker's costs). Individuals are essentially betting against one another in a situation where the odds change over the course of the betting period. Thus, they do not know the odds they will receive if their horse wins at the time they place their bets. Both parimutuel and financial markets are therefore characterized by uncertainty about future payoffs as well as by investors who have potential access to some private information, the history of trades, and a statistic which reflects the market information on winning probabilities. This statistic consists of the current odds in betting markets or the price in financial markets. Unlike in financial markets, though, the uncertainty is resolved unambiguously at the time the outcome is realized in betting markets. This facilitates the precise measurement of market informational efficiency and its comparison to benchmark levels.

As a consequence, numerous field studies on horse race betting have attempted to statistically test the extent to which betting markets are efficient (see Vaughan Williams, 2005, chap. 2 and 3 for extensive reviews). Starting with Griffith (1949), a number of empirical studies have used racetrack betting data to examine the relationship between the objective probability of winning a bet and its market (final) odds. By investigating the expected returns to bets placed at differing odds, these studies have concluded 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 odds reflect a tendency of bettors to underbet on favorites and to overbet on longshots. The implication of this bias is that bettors can make above-average (though not usually profitable) returns by betting on favorites, which, in turn, is evidence of weak form information inefficiency in betting markets. Studies which have investigated the existence of strong form efficiency in racetrack betting markets have addressed the issue of the value of inside information in these markets (Schnytzer and Shilony, 1995). The evidence suggests that insiders in betting markets possess valuable information unavailable to the public, which they can trade upon so as to earn above-average and even abnormal returns, and to this extent the market may be considered informationally inefficient. Note that, since the information that is held by the bettors is unknown to the researchers, field tests have used the absence of opportunities to earn abnormal returns to determine whether the market is efficient. In contrast, incentives and information can be carefully controlled in laboratory markets. Laboratory studies therefore provide a unique opportunity to *directly* address the question of information efficiency in betting markets.

The experiment reported in this paper concerns market pricing and information aggregation in

a parimutuel betting market. The primary focus of the paper is to consider whether elicitation of beliefs from individuals, about the probability that each outcome will occur, improves the information aggregation capacity of the market. There is reason to believe that belief elicitation might exert a considerable effect on betting decisions and, as a consequence, on prices. Requiring bettors to state beliefs may actually influence the beliefs themselves, and thus bettors' trading strategies. They may also, however, change the weight the outcome probabilities receive in the bettor's decision process, relative to the odds, and thereby affect bettor behavior. These effects may encourage better decisions, and thus improve the information aggregation performance of the market. In particular, belief elicitation may reduce the tendency to bet on longshots by emphasizing the relatively low probability that they will occur, and thus mitigate the favorite-longshot bias widely observed in previous studies. There is also evidence, from other studies, that belief elicitation encourages more rational behavior in other contexts.¹

Plott, Witt, and Yang (2003) present the first experimental evidence on the capacity of parimutuel betting markets to aggregate information. They study a setting in which asymmetrically informed individuals could place bets on outcomes by making purchases from a market maker in several markets operating 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, and that a small favorite-longshot bias exists. Due to the complexity of the environments that were implemented, this analysis mainly focuses on market prices and does not analyze the betting decision process. On the contrary, we study a simpler betting environment where we mainly address the role of belief elicitation on betting decisions. The structure of our market is related to that modeled in Avery and Zemsky (1998) and studied experimentally in Drehmann, Oechssler, and Roeder (2005) and in Cipriani and Guarino (2005).² These two experimental studies found widespread contrarian behavior, which is suboptimal betting on the longshot and against one's private information. This contrarian betting distorted market prices and generated a favorite-longshot bias. In our study, we measure the change in the incidence of contrarian behavior resulting from the elicitation of beliefs.³

¹The previous experimental literature on belief elicitation is primarily concerned with strategic uncertainty, and measuring expectations about others' strategies in normal form games. Most of the experimental evidence suggests that prompting subjects for beliefs about others' strategies moves their choices closer to equilibrium (see, among others, Croson, 2000).

²This structure is a straightforward extension of an experimental design used to study information cascades (Anderson and Holt, 1997), with the additional feature that odds on each outcome exist and are updated after each individual makes his bet.

³Drehmann, Oechssler, and Roeder (2005) conducted an experimental test of Avery and Zemsky's (1998) model. As in the model, the price is automatically set equal to the level that would result if the market maker infers the private information of bettors from their bets, when such inference is possible, and prices each outcome fairly given this information. Thus, it is always optimal for each bettor to bet in accordance with his private information. This means that using the simple heuristic of following one's own signal leads to optimal behavior and drawing inferences from the sequencing of prior bets is not necessary. Nevertheless, only 50 - 66 percent of agents follow their own signal, and there are considerable differences between observed odds and objective probabilities of the outcomes. There is a high incidence of contrarian behavior (betting on the current long-shot as well as against one's own private information), which is never optimal in their setting. Cipriani and Guarino (2005) study a similar sequential betting game. They consider three treatments that differ in the pricing rule that the market maker uses and in the information available to bettors. In one treatment, prices are independent of betting history. In the other two treatments, prices are set as in the Drehmann, Oechssler, and Roeder study, but the treatments differ in the availability of the prior betting history. They find a considerable amount of suboptimal behavior, consisting mainly of contrarian betting.

Our experiment consists of three different treatments, which are described in detail in section two. In all treatments, subjects observe private signals about the probability that each of two a priori equally likely events has occurred. They then each sequentially make a bet on one of the two possible outcomes, and a winning bet yields a return that is decreasing with the proportion of subjects who have bet on the winning outcome. In the *Bet* treatment, subjects make bets only. In the second treatment, *ObsPred*, there are two groups of subjects: bettors and observers. Bettors make bets while observers submit beliefs about the probabilities of each of the possible final outcomes based on the betting data they have observed. In the third treatment, *BetPred*, the same subjects submit both bets and beliefs about the final outcomes. In the last two treatments, belief statements are not made public. That is, no subject observes at any time the beliefs any other subject has submitted.

We compare the three treatments with regard to the extent that the final market odds reflect information aggregation. Comparing market behavior between treatments *Bet* and *BetPred* allows us to consider whether requiring bettors to submit beliefs leads to more successful information aggregation. Comparison of belief statements in treatments *ObsPred* and *BetPred* allows us to investigate whether making bets affects the accuracy of stated beliefs, that is, whether the act of betting influences belief formation. Comparing bets in treatments *Bet* and *ObsPred* reveals whether bettors' decisions change when others observe them. That is, does making it common knowledge that beliefs are being elicited, with the possible corresponding change in emphasis on outcome probabilities, affect betting behavior? Our results are reported in section three. We find that the market's capacity to aggregate information is better when bettors are required to submit beliefs, in *BetPred*, than in the other two treatments. The favorite-longshot bias, a phenomenon observed in treatments *Bet* and *ObsPred*, is reduced when beliefs are elicited from bettors in *BetPred*. Thus, eliciting beliefs improves betting decisions and, as the comparison between the *ObsPred* and *BetPred* treatments reveals, placing bets improves the accuracy of belief statements.

We then turn to individual behavior in detail in section four. To attempt to explain the individual betting data, we propose a 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. The Adaptive Model is appealing as a descriptive model because, unlike standard game-theoretic equilibrium concepts, it (1) generically makes a unique prediction at every decision node, and (2) agents make inferences only from past history. We argue that the Adaptive Model performs very well in describing behavior in the *BetPred* treatment. We advance a conjecture that the reason that submission of beliefs improves the performance of the market is because it encourages individuals to distinguish between informative and uninformative prior bets when making their betting decision.

2 The Experiment and Measures of Information Aggregation

2.1 The Environment

Consider a “horse” race with two horses called A and B . There is a finite set $N \equiv \{1, \dots, n\}$ 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 i th bettor in the sequence. Each bettor observes the betting decisions of all previous individuals 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 $s = (s_1, \dots, s_n) \in \{A, B\}^n$ 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 \bar{H} denote the horse other than H . The *odds against horse H* , given by the total number of bets on horse \bar{H} divided by the total number of bets on horse H , is denoted by

$$O_H(s) = \frac{n - h(s)}{h(s)}. \quad (1)$$

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.

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

$$\begin{aligned} \Pr(q_i = q^A \mid \theta_A, q_j) &= \Pr(q_i = q^A \mid \theta_A) = \pi \in (1/2, 1) \\ \Pr(q_i = q^A \mid \theta_B, q_j) &= \Pr(q_i = q^A \mid \theta_B) = 1 - \pi, \end{aligned} \quad (2)$$

for all $i, j \in N$, $i \neq j$. Hence, bettor i ’s interim beliefs⁴ are given by $\Pr(\theta_H \mid q^H) = \pi$ and $\Pr(\theta_H \mid q^{\bar{H}}) = (1 - \pi)$, $H = A, B$.

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, placed bets and stated beliefs in each of the three treatments.

⁴Bettors’ beliefs at the interim stage correspond to the beliefs bettors entertain *after* they have received their private information but *before* they have observed any market activity.

Table 1: NUMBER OF SESSIONS, GROUPS, SUBJECTS, BETS, AND BELIEF STATEMENTS IN EACH TREATMENT

	<i>Bet</i> (subjects bet only; 20 rounds)	<i>ObsPred</i> (1/2 of subjects bet, 1/2 state beliefs; 20 rounds)	<i>BetPred</i> (same subjects bet and state beliefs; 20 rounds)
Sessions	2	4	2
Groups	4	8 (4 placing bets, 4 stating beliefs)	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 the *Bet* and *BetPred* treatments, and one of the two groups in the *ObsPred* treatment, played 20 repetitions of the game 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 game, while each *period* refers to one subject’s turn to bet within a round. Subjects were instructed on the rules of the game and the use of the computer program⁵ 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. Afterwards, the twenty rounds of the sequential betting game that constituted the experiment took place.⁶ Communication between the subjects was not allowed. Each session was between 90 and 135 minutes in duration.

2.2.1 The Betting Process

Each round involved the following sequence of events. At the beginning of a round, a random choice was made between color *A* (state θ_A) and color *B* (state θ_B), with a probability of choosing color *A* equal to 1/2. 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.⁷ 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

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

⁶Twenty-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.

⁷The 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).

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.

In all sessions, the same random draws were used, so that the same signals were assigned to the same positions in the sequence of bettors in each round. However, the draws differed between rounds within a given session. Thus, in a given round t , the same nine random draws (one state of Nature and eight private signals conditional on the state of Nature) were used for each group of participants, though the draws in round $s \neq t$ differed from those in round t . At the end of each session, subjects received six euros for each ECU they won during a subset of the twenty rounds played.⁸

2.2.2 Belief Elicitation

The group that did not make bets in *ObsPred* 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, during the current round, 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 reported 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 .⁹

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 reported belief vector was given by $\text{€ } .15 \left(1 - (\mu_i^k)^2 \right)$ 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 prediction, 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.¹⁰

In the *BetPred* treatment, each group of eight participants played twenty repetitions of the betting game. Each participant, in addition to making bets in his designated period, reported his beliefs in each period about the likelihood that each color was the true state, in an identical manner to *ObsPred*. In *BetPred*, participants observed only the betting activity of their own group.¹¹

⁸At 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 of whether the random lottery incentive procedure distorts behavior, find that random lottery designs yield behavior no different from one-shot games.

⁹In the experiment, subjects entered μ_i^A and μ_i^B as numbers in the interval $[0, 100]$, which were described as percentages.

¹⁰Belief elicitation using a quadratic scoring rule is widely employed in experimental economics (see for example Nyarko and Schotter, 2002). Offerman, Sonnemans, van de Kuilen, and Wakker (2007) show how proper scoring rules can be generalized to modern theories of risk and ambiguity, and can become valid under risk aversion and other deviations from expected value. They also report experimental results suggesting that it is desirable to correct subjects’ reported probabilities elicited with scoring rules if only a single large decision is paid but that this correction is unnecessary with many repeated decisions and repeated small payments.

¹¹In the *Bet* and *ObsPred* treatments, three of the twenty rounds were chosen at random at the end of the session to count toward bettors’ earnings. In *BetPred*, 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 both *ObsPred* and *BetPred*.

2.3 Measuring Information Aggregation

In this subsection, we define the criterion we use to compare the level of information aggregation between the treatments. Consider the decision of bettor i , the i th bettor in the sequence. Define the history of bets up to and including bettor i 's bet as $s^i = (s_1, \dots, s_i) \in \{A, B\}^i$. s^i implies that the odds in period i against horse H are given by

$$O_H(s^i) = \frac{i - h(s^i)}{h(s^i)}. \quad (3)$$

The odds after period i define a *market probability* that horse H wins the race. This market probability, P_H^m , is given by

$$P_H^m(s^i) \equiv \frac{h(s^i)}{i} = \frac{1}{O_H(s^i) + 1}. \quad (4)$$

After all bets have been made, i.e., $s^i = s$, this market 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. While the market probability is based on the profile of bets, there are several different interesting benchmark price vectors, corresponding to different probabilities. One such benchmark, the *objective probability* after period i that horse H wins the race, equals

$$P_H^O(q^i) \equiv \Pr(\theta_H | q^i), \quad (5)$$

where $q^i = (q_1, \dots, q_i)$ 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 . If prices reflect the objective probabilities, betting on each outcome offers an equal expected return.

To measure the extent to which information aggregation takes place in our markets we use two other baselines. The first baseline is the price in an hypothetical parimutuel betting market, in which each bettor has access to all of the private information of previous bettors, as well as his own. This baseline represents the maximum possible level of information aggregation the market could achieve. It differs from the objective probability because it reflects odds established with bets made with the information known to the market at the time each bettor placed his bet. The dynamics of prices in such a market are straightforward to characterize: bettor i bets on horse H if $\Pr(\theta_H | q^i) > \frac{h(s^{i-1})+1}{i+1}$. The resulting final *observable signals* (OS) market 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. In such a market the market probability of each outcome equals on average 1/2 no matter what the signal profile. Now define

$$V = \sum_{q^8 \in Q_E^8} \frac{1}{20} \left| P_A^{OS}(q^8) - P_A(q^8) \right|, \quad (6)$$

and

$$V_{\max} = \sum_{q^8 \in Q_E^8} \frac{1}{20} \left| P_A^{OS}(q^8) - 1/2 \right| \quad (7)$$

where q^8 denotes a signal profile ($n = 8$), $P_A(q^8)$ denotes the observed market probability for horse A after period 8, and Q_E^8 denotes the set of twenty signal 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.

3 Results


Figure 1 plots the time path of horse A 's objective and market probability for each round of the experiment. A separate time series is provided for each of the three treatments for each round, and the data are averaged over the four groups 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 market probability typically implies the same favorite as the objective probability. The only exceptions are in rounds 6 and 12 for *Bet*, and rounds 6, 11, and 19 for the *ObsPred* and *BetPred* treatments. The final market probability in the *BetPred* treatment is usually (in 15 of 20 rounds) closer to the objective probability than in each of the other treatments, suggesting better information aggregation in *BetPred*. There does not seem to be a systematic difference with respect to information aggregation between the *Bet* and *ObsPred* treatments, as average final market probabilities in *ObsPred* are closer to objective probabilities than in *Bet* for 12 of 20 rounds. Deviations of the market 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.

Result 1 provides a clear statement of our results with regard to information aggregation by treatment. The statistical support for the result is provided in the paragraph that follows.

Result 1 (Information aggregation): Elicitation of beliefs of bettors increases the information aggregated in the market.

Support: The values of IA observed in treatments *Bet* and *ObsPred*, 50.45% and 53.18%, are lower than the 68.18% in *BetPred*. The average absolute differences between the market and objective probabilities are .180, .174, and .138 in treatments *Bet*, *ObsPred*, and *BetPred*, respectively, for the complete data for all eight periods that make up each round. For the data from the last period only, the absolute differences are .299, .280, and .210 in the *Bet*, *ObsPred*, and *BetPred* treatments, respectively. The average value of IA for each of the four groups of bettors that participated in *BetPred* is higher than the average values of IA for any of the four groups in either *Bet* or *ObsPred*. Thus, a conservative rank-sum test, using each group of participants as an observation, rejects the hypotheses

Figure 1: OBJECTIVE AND MARKET PROBABILITIES IN EACH TREATMENT, AVERAGED ACROSS GROUPS



at the 5 percent level that (a) treatments *Bet* and *BetPred*, and (b) treatments *ObsPred* and *BetPred*, aggregate the same amount of information. A similar test for a difference between treatments *Bet* and *ObsPred* is insignificant at conventional levels. □

A comparison between market and objective probabilities implied by the final odds at the end of period eight shows that the market probability on the favorite is usually lower than the objective probability. This indicates that a fictitious bettor, placing a bet at the final odds, would have a higher expected payoff betting on the favorite than on the longshot. The extent of this favorite-longshot bias, and the difference between treatments, is summarized in result 2.

Result 2 (Favorite-longshot bias): Market odds are consistent with a favorite-longshot bias in the *Bet* and *ObsPred* treatments. The bias is reduced when bettors' beliefs are elicited in the *BetPred* treatment.

Support: By the end of the round, the objective probability for one of the horses equals at least 0.9 in every round. The final market probabilities are always divisible by $1/8$, so that market probabilities equal to 1 or $7/8$ in favor of the objective favorite do not necessarily indicate a deviation from the objective probability, but any market probability of .75 or lower on the favorite is consistent with a favorite-longshot bias. This occurs in 87.50% of markets in the *Bet* treatment, and 68.75% in *ObsPred*, but only 43.75% in *BetPred*.¹² \square

Figure 2 plots the average¹³ time path of horse *A*'s objective probability compared to the average stated belief in *ObsPred* and *BetPred*. It is evident from the figure that the differences between the two treatments are considerable. Stated beliefs at the end of period 8 in *BetPred* are on average closer to the objective probability than in *ObsPred* in all 20 rounds. Average beliefs in *ObsPred* are closer to $1/2$, an equal probability on each outcome, than in *BetPred* in 17 of 20 rounds. The exceptions are rounds 6, 11, and 19, the only ones in which the final market probabilities yield an incorrect prediction in a majority of treatments.

As we report in result 3 below, we also observe that when the beliefs are submitted by the same individuals who make the bets, the beliefs are closer to reflecting the market probabilities. This may not be surprising since when predictors and bettors are different individuals, as in *ObsPred*, they might have different opinions. However, as result 4 shows, the beliefs are also closer to the objective probabilities when the same individuals submit beliefs and make bets, indicating that the act of betting improves the accuracy of predictions.

Result 3 (Consistency between betting data and stated beliefs): Submitting bets in addition to beliefs improves the consistency between betting data and stated beliefs.

Support: We measure the degree of consistency between the betting data and the stated beliefs using the average absolute difference between an individual's stated beliefs and the market probabilities. All four groups in *BetPred* have smaller average absolute differences than any group in *ObsPred*. 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 stated beliefs and market probabilities in *ObsPred* and *BetPred* are equal. \square


Result 4 (Effect of placing bets on accuracy of beliefs): Submitting bets improves the accuracy of belief statements.

Support: Table 2 shows the average value of the absolute difference between the average stated belief and the objective probability in *ObsPred* and *BetPred* in each group. Recall that each group participated in exactly one treatment, so that the data in the row of Table 2, corresponding to a given group, pertains to different individuals in each treatment. The average difference for each of the four groups of bettors that participated in *BetPred* is lower than the average absolute differences

¹²We recognize that in many parimutuel betting markets, including ours, pure randomness in decisions would also tend to produce a favorite-longshot bias.

¹³Because 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 stated belief and the objective probability of horse *A*, we compute the average objective probability in period *i* as $(7 \Pr(\theta_A | q^{i-1}) + \Pr(\theta_A | q^i)) / 8$.

Figure 2: OBJECTIVE PROBABILITY AND STATED BELIEFS IN *ObsPred* and *BetPred*, AVERAGED ACROSS GROUPS



for any group in *ObsPred*, whether all periods or only the final periods are used in the calculation. In addition, subjects' stated beliefs in *BetPred* lead to earnings from predictions that are 14.68% higher than those in *ObsPred*. If the comparison between treatments is restricted to instances in which the previous sequences of bets and private signals are identical, the difference in earnings is 10.46%.

□

4 Individual Behavior

In this section, we propose a model, called the *Adaptive Model*, and consider whether it can explain betting decisions in our experiment. The sequential parimutuel betting game we study in our experiment is a well-defined extensive form game, and in principle, Nash or sequential equilibria can be

Table 2: AVERAGE ABSOLUTE DIFFERENCE, IN TERMS OF PROBABILITY OF OUTCOME, BETWEEN BELIEF STATEMENTS AND OBJECTIVE PROBABILITY, EACH GROUP, ALL ROUNDS

	<i>ObsPred</i>		<i>BetPred</i>	
	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
Average for all four groups	.2022	.3025	.0878	.0908

identified, although this is an intractable exercise due to the enormous potential number of strategy profiles (see Koessler, Noussair, and Ziegelmeyer, 2007). However, there are two difficulties with using sequential equilibrium or any weaker equilibrium concept as a predictive tool for this game. 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 generic existence of multiple equilibria means that the predictions that can be extracted from equilibrium analysis are typically indeterminate.

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.¹⁴ Rather, 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. They update their beliefs about the outcome probabilities 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 number of experimental studies suggesting that expectations are adaptive rather than rational.¹⁵ The model generically yields a unique predicted action at each decision node.

Suppose a bettor perceives the payoff of betting on H in period i as

$$U_i(s^i, \theta) = \begin{cases} O_H(s^i) + 1 = \frac{i}{h(s^i)} & \text{if } s_i = H \text{ and } \theta = \theta_H, \\ 0 & \text{if } s_i = H \text{ and } \theta \neq \theta_H. \end{cases} \quad (8)$$

¹⁴Drehmann, Oechssler, and Roeder (2005) elicit predictions of market prices one period in advance. They note that predictions typically fail to take into account the effect of the current bet on future prices. Their method of eliciting predictions is different from ours because we elicit predictions about the objective probability, rather than future prices, which correspond to future market probabilities. It stands to reason, therefore, that taking into account the effect of one’s current bet on future betting behavior is an assumption that is likely to be inappropriate for a behavioral model.

¹⁵For example, Marimon and Sunder (1993), Camera, Noussair, and Tucker (2003) and Johnson, Camerer, Sen, and Rymon (2002) find greater empirical support for adaptive rather than rational expectations in three different contexts. Ziegelmeyer, Broihanne, and Koessler (2004) find that the betting behavior in their parimutuel betting markets is inconsistent with backward induction.

Let k and \bar{k} denote the number of signals favoring horse H and \bar{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 $k + \bar{k} < i - 1$ is possible. Accordingly, bettor i 's posterior probability that horse H wins the race is given by

$$\mu_i(\theta_H | s^{i-1}; q_i) = \begin{cases} \left(\pi^{k+1} (1 - \pi)^{\bar{k}} \right) / \left(\pi^{k+1} (1 - \pi)^{\bar{k}} + (1 - \pi)^{k+1} \pi^{\bar{k}} \right) & \text{if } q_i = q^H, \\ \left(\pi^k (1 - \pi)^{\bar{k}+1} \right) / \left(\pi^k (1 - \pi)^{\bar{k}+1} + (1 - \pi)^k \pi^{\bar{k}+1} \right) & \text{if } q_i = q^{\bar{H}}. \end{cases} \quad (9)$$

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 \bar{H} . This is equivalent to saying that bettor i bets on horse H if

$$\mu_i(\theta_H | s^{i-1}; q_i) > \frac{h(s^{i-1}) + 1}{i + 1}. \quad (10)$$

Equation (10) 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 $q^{\bar{H}}$ leads to $s'_k \neq s_k$. Second, bettor i forms his posterior probability that horse H wins the race using equation (9). Finally, he bets on horse H if inequality (10) 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.

As an alternative model against which to evaluate the Adaptive Model, we propose a plausible betting rule called the *Private Information Plus Odds Heuristic* (INFODDS). A bettor using this heuristic does not infer any information from the history of past bets but simply responds optimally to the current odds given his private signal. Formally, bettor i bets on horse H if $\Pr(\theta_H | q_i) > \frac{h(s^{i-1})+1}{i+1}$. For example, when his private signal is q^A and $\pi = 3/4$, he bets on horse A unless the return from horse B is more than three times the return on horse A . The difference between INFODDS and the Adaptive Model is that the Adaptive Model takes into account the information embodied in the sequencing of prior bets. In the Adaptive Model, the bettor infers the number of signals of each type held by previous bettors, which requires the individual to distinguish between informative and uninformative bets.

We consider the extent to which individual bets are consistent with the Adaptive Model and INFODDS, and how the level of consistency varies by treatment. Our results are summarized below.

Result 5 (Betting behavior, the Adaptive Model and InfOdds): Individual bets in all treatments provide strong support for the Adaptive Model. This is particularly true in the *BetPred* treatment. In all treatments, the Adaptive Model correctly predicts more individual betting decisions than InfOdds. The difference is greatest in the *BetPred* treatment.

Support: We compare actual betting decisions with the predictions of the Adaptive Model and of INFODDS.¹⁶ The predictions of these two benchmarks coincide in most instances, but in 23.91% of the observations in the *Bet* treatment, 26.09% in *ObsPred*, and 33.59% in *BetPred*, the Adaptive Model and INFODDS make different predictions. The second row of table 3 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 *BetPred* than in *Bet* and *ObsPred*, between which the difference is small. The Adaptive Model is therefore more relevant for predicting decisions made when bettors explicitly state their beliefs.

The third row of data in table 3 indicates the percentage of bets consistent with INFODDS. Comparison of rows two and three shows that more bets are consistent with the Adaptive Model than with INFODDS in all three treatments. The fourth row gives the percentage of bets consistent *only* with the predictions of the Adaptive Model in those instances in which it and INFODDS make different predictions. In all three treatments a significant majority of the participants' decisions in these cases 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 INFODDS in all three treatments, and the difference is especially pronounced in *BetPred*.

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 INFODDS. The fact that, in each treatment, approximately half of the bets inconsistent with the Adaptive Model are also inconsistent with INFODDS indicates that INFODDS is of no additional value as a predictor over and above the Adaptive Model (since predicting bets in accordance or alternatively counter to INFODDS for observations counter to the Adaptive model yields essentially the same hit rate). □

Result 5 establishes that requiring participants to state their beliefs leads to bets that are closer to the Adaptive Model predictions. One possible explanation for this 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. This could occur, perhaps, because the placement of bets without belief elicitation may cause an overemphasis on the payoff from each outcome given the odds, rather than on its probability of occurrence. Result 6 indicates that belief elicitation does tend to reduce particular types of deviations more than others.

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 information¹⁷; ii) *Incorrect Herding*: betting on the favorite and against one's own private information when doing so is inconsistent with the Adaptive

¹⁶If the previous sequence of decisions is inconsistent with any possible sequence for the Adaptive Model then we assume that it is common knowledge that any bettor, who has deviated from the Adaptive Model, followed 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.

¹⁷We include betting against one's private information when the current odds on each outcome are equal, in the category of contrarian betting.

Table 3: PERCENTAGE OF DECISIONS CONSISTENT WITH THE ADAPTIVE MODEL AND INFODDS

	<i>Bet</i>	<i>ObsPred</i>	<i>BetPred</i>
Consistent with the Adaptive Model	81.09% (519/640)	81.25% (520/640)	87.03% (557/640)
Consistent with INFODDS	75.62% (487/640)	68.28% (436/640)	65.94% (422/640)
Consistent with the Adaptive Model in instances in which the Adaptive Model and INFODDS make different predictions	61.44% (94/153)	74.85% (125/167)	81.40% (175/215)
Neither consistent with the Adaptive Model nor with INFODDS	9.69% (62/640)	12.19% (78/640)	6.72% (43/640)

Model; and iii) *Failure to Herd*: betting according to one’s own private information and against the favorite when doing so is inconsistent with the Adaptive Model. Note that betting on the favorite when it coincides with one’s own private signal is always consistent with the Adaptive Model.

Result 6 (Effect of belief elicitation on deviations from the Adaptive Model): Elicitation of beliefs from bettors 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: Table 4 reports, for each treatment, the percentage of instances in which each type of deviation occurred when it was possible. Contrarian behavior and failure to herd are less frequent in *BetPred* than in the other two treatments. Contrarian behavior decreases by 71% from the level in *Bet* and by 73% from the level in *ObsPred*. On the other hand, the incidence of incorrect herding is higher in *BetPred* than in the other two treatments.

Table 4: PERCENTAGE OF OCCURRENCE OF THE THREE TYPES OF DEVIATION FROM THE ADAPTIVE MODEL (OF INSTANCES IN WHICH DEVIATION WAS POSSIBLE)

	<i>Bet</i>	<i>ObsPred</i>	<i>BetPred</i>
Contrarian Behavior	10.55% (46/436)	11.49% (51/444)	3.10% (14/452)
Incorrect Herding	22.13% (27/122)	29.52% (31/105)	34.38% (30/87)
Failure to Herd	58.54% (48/82)	41.76% (38/91)	38.61% (39/101)

□

We now show that the differences in the accuracy of the Adaptive Model across treatments, and the relative incidence of the different types of deviations from the model, are consistent with a statistical

generalization of the model with less costly deviations in *BetPred* than in the other two treatments. Under the Adaptive Model, the perceived expected utility to individual i from betting on horse H is $EU_i(H | s^{i-1}; q_i) = \mu_i(\theta_H | s^{i-1}; q_i) \frac{i}{1+h(s^{i-1})}$ where $\mu_i(\theta_H | s^{i-1}; q_i)$ is bettor i 's belief that horse H will win the race given the information available. Suppose that bettor i bets on horse H with probability

$$\Pr(i \text{ bets on } H) = \frac{e^{\lambda EU_i(H|s^{i-1};q_i)}}{e^{\lambda EU_i(A|s^{i-1};q_i)} + e^{\lambda EU_i(B|s^{i-1};q_i)}}, \quad (11)$$

where “errors,” i.e., bets on the horse with the lower expected payoff under the Adaptive Model, are measured in terms of a precision parameter $\lambda \geq 0$. If $\lambda = 0$, then individuals are betting completely randomly, while for $\lambda = \infty$, individuals’ bets perfectly match the Adaptive Model’s predictions. For other positive values of λ , individual i is more likely to bet on horse H , the greater the expected payoff from doing so compared to betting on the other horse, but with a probability less than one. 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 (McKelvey and Palfrey, 1998). Our findings are summarized in the statement of the next result.

Result 7 (QRAM): Assuming that betting decisions are characterized by QRAM, the elicitation of beliefs from bettors reduces the estimated incidence of costly deviations from the Adaptive Model as measured by the parameter λ .

Support: Table 5 reports the results of the maximum-likelihood estimation of QRAM, for the three treatments. The table also contains the standard error of each estimated parameter in parentheses. The negative log likelihood is given in the last row. Likelihood ratio tests indicate that the precision parameter λ is not significantly different between *Bet* and *ObsPred*, but it is significantly greater in *BetPred* than in each of the other two treatments. This higher precision is consistent with better predictive power of the Adaptive Model in *BetPred* than in the other two treatments.

Table 5: ESTIMATED PRECISION PARAMETER OF QRAM FOR EACH OF THE THREE TREATMENTS

	<i>Bet</i>	<i>ObsPred</i>	<i>BetPred</i>
λ	2.890 (.193)	2.837 (.257)	7.184 (.590)
$-ll^*$	277.288	301.725	187.994

Notes: Numbers in parentheses are the standard error of the parameter. $-ll^*$ is the negative log likelihood.

□

To verify whether the difference in the precision parameter between treatments is consistent with the pattern of deviations from the Adaptive Model reported in table 4, we generated sequences of bets according to QRAM for numerous values of λ and then we computed the relative frequencies of each type of deviation. The results show that contrarian behavior and failure to herd decline

monotonically as λ increases (contrarian behavior declines very rapidly). However, there is a range of moderate λ values (between approximately 2 and 10) for which the incidence of incorrect herding deviations *increases* as the error rate *decreases*. Therefore, the joint hypothesis of QRAM and less costly deviations from the Adaptive Model in *BetPred* than in the other two treatments explains the treatment differences for all three types of error: the lower incidence of contrarian betting and failure to herd as well as the higher incidence of incorrect herding.

5 Conclusion

We find that pricing patterns reflect greater information aggregation in *BetPred*, in which bettors themselves must state beliefs, than in the other two treatments. Experimental economists have found that eliciting beliefs of subjects is a useful methodological tool for theory testing or 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 predict behavior. Our results suggest yet another use for belief elicitation procedures, as potential instruments to improve a market's capacity to aggregate information.

This finding is of special interest for a new form of financial market, often known as a prediction market. Prediction markets are exchange institutions designed with the sole purpose of forecasting future events like election outcomes, corporate sales or economic indicators (Wolfers and Zitzewitz, 2004). These markets have in many cases been shown to forecast future events better than experts or opinion polls, with the exception of low-probability events. Our conjecture is that eliciting traders' beliefs about future market probabilities would improve the potential of prediction markets to aggregate and organize information that is held in a widely dispersed and subjective form.

The *BetPred* treatment is characterized by less frequent failure to herd when herding is arguably 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, *BetPred* exhibits a much smaller favorite-longshot bias than the other two treatments. It appears that the attractive payouts for longshots in the event that they win, coupled with a poor estimate of the probability that the longshot is victorious, leads to decision errors. However, the elicitation of beliefs appears to improve bettors' probability assessments, which in turn discourages an overemphasis on the attractive return associated with a longshot during the decision process.

We proposed a model, called the Adaptive Model, to characterize betting decisions. It assumes that individuals distinguish informative from uninformative previous bets, connect informative bets to underlying signals, calculate the probability of each outcome implied by multiple imprecise signals, and compare the result to the current odds. On the other hand, the Adaptive Model assumes that bettors do not take into account the effect of their own bet on the behavior of later bettors. The model is conducive to empirical testing because it makes a unique prediction at each node, rendering it potentially falsifiable.

The data suggest that the rationality assumptions of the Adaptive Model are reasonable. The model is successful despite the fact that the betting decisions are made in an environment where

the current odds, which give equal weight to informative and uninformative bets, may at times be misleading about informational content, and where individuals at times must override their own private information. The superior performance of the Adaptive Model over the Private Information Plus Odds Heuristic indicates that agents do infer and use the informational content of prior bets, which requires the proper use of sequencing information, in making their own betting decisions. It appears that eliciting beliefs from bettors facilitates this inference about the number of underlying signals based on the position of previous bets in the betting sequence, and thus helps them distinguish between informative and uninformative bets.

Betting decisions are no better when they are observed than when they are not observed, indicating that common knowledge of the fact that beliefs are being elicited is insufficient to improve decision-making. *Bettors* must be submitting their *own* beliefs to lead to improved betting decisions. Stated beliefs are also more accurate, as well as more consistent with market activity, when the agents submitting beliefs are also placing bets. On the other hand, when beliefs are submitted by observers rather than bettors, they tend to anticipate more noise in the decision making of bettors than the decisions actually exhibit. This may be the case because it is harder 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 case the bettors, relative to their own level.

References

- ANDERSON, L. R., AND C. A. HOLT (1997): "Information Cascades in the Laboratory," *American Economic Review*, 87(5), 847–862.
- AVERY, C., AND P. ZEMSKY (1998): "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88, 724–748.
- BOUN MY, K. (2002): "A VB Application for Sequential Betting with Asymmetric Information," CNRS, BETA-Theme.
- CAMERA, G., C. NOUSSAIR, AND S. TUCKER (2003): "Rate of Return Dominance and Efficiency in an Experimental Economy," *Economic Theory*, 22(3), 629–660.
- CIPRIANI, M., AND A. GUARINO (2005): "Herd Behavior in a Laboratory Financial Market," *American Economic Review*, 95(5), 1427–1443.
- 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 (2005): "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment," *American Economic Review*, 95(5), 1403–1426.

- GRIFFITH, R. M. (1949): “Odds Adjustments by American Horse-Race Bettors,” *American Journal of Psychology*, 62, 290–294.
- 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.
- KOESSLER, F., C. NOUSSAIR, AND A. ZIEGELMEYER (2007): “Parimutuel Betting under Asymmetric Information,” *Journal of Mathematical Economics*, In Press.
- 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. 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, G. VAN DE KUILEN, AND P. P. WAKKER (2007): “A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes,” Econometric Institute, Erasmus University, Rotterdam, The Netherlands.
- 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.
- 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.
- 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.
- (ed.) (2005): *Information Efficiency in Financial and Betting Markets*. Cambridge University Press.

WOLFERS, J., AND E. ZITZEWITZ (2004): "Prediction Markets," *Journal of Economic Perspectives*, 18(2), 107–126.

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

Information Aggregation and Beliefs in Experimental Parimutuel Betting Markets: Instructions

(included only for the referees' inspection)

The original instructions were written in French. Here we include only the translation of the instructions used in treatment three, in which participants both place bets and state beliefs about the likelihood of each outcome. The instructions for the other two treatments involve only minor changes from those reprinted here.

Instructions

WELCOME

This is an experiment about interdependent decision making. The instructions are simple and if you follow them and make good decisions, you may earn a considerable amount of money. Your earnings will depend partly on your decisions and partly on chance. All of your decisions will be treated in an anonymous manner and they will be gathered across a computer network. You will input your choices on the computer you are seated in front of and the computer will indicate your earnings to you as the experiment proceeds. The total amount of money that you earn in the experiment will be given to you in cash at the end of the experiment.

Before beginning the experiment, you must fill out a questionnaire intended to evaluate your comprehension of the instructions. If your answers indicate that you have not understood the following instructions well enough, you will not be able to take part in the experiment. If this occurs, we will ask you to leave and you will receive 5 euros at the exit.

One of you will be designated at random to assist the experimenters during the session. This person will be referred to as the assistant in the remainder of these instructions. The assistant will read the instructions aloud, will observe how the session is proceeding to assure you that we are respecting the instructions, and, as described later, will draw some random numbers. The assistant will receive a final payment equal to the average amount paid to the participants.

GENERAL SETTING OF THE EXPERIMENT

16 people will participate in the experiment. Two groups of eight people will be formed at random at the beginning of the experiment. There will be no way for you to identify which of the other participants are in your group, because they can be seated anywhere in the room. At the beginning of the experiment you will receive 20 tokens.

The experiment will consist of **20 markets**. During each market, you will place one token. **Each market** is divided into **8 periods**. In each period, one of you will be required to make a decision. The possible decisions you can make are to “place 1 token on color X ” or “place 1 token on color Y .” Each of you will place exactly one token in each market. In each market, the period during which you will make your choice will be determined at random. Your decision can translate into earnings at the end of each market. Furthermore, over the course of each market, you will be required to make a certain number of predictions. Your predictions will lead to *bonus earnings* at the end of each market. The earnings from your placement of tokens and from

your predictions are denominated in tokens and at the end of the experiment the tokens will be converted into euros. The conversion procedure will be explained at the end of the instructions. The remainder of the instructions will allow you to understand how your earnings will be determined.

HOW THE MARKET OPERATES

Each market begins with a random draw of a winning color from the two possible colors. The winning color drawn at random will be either the color X or the color Y , where the actual colors called X and Y will change from one market to the next. The assistant will conduct the random draw of the winning color. At the beginning of each market, each color has the same chance of being drawn, i.e., the chance that color X will be chosen is 50% and the chance that color Y will be chosen is 50%. For each market, you will not know which color the assistant drew until the end of the market. After the random draw takes place for a market, period 1 begins.

In each period, only one person will place a token on a color. The participant assigned to make a decision for the current period must choose between “place 1 token on color X ” and “place 1 token on color Y .” The *end of period* screen will appear on all participants’ computer screens as soon as the person assigned to the current period has made his or her decision. The end of period screen will present the **odds** of each of the two colors. The odds are calculated in the following manner:

$$\text{Odds on color } X = \frac{\text{Number of tokens placed on color } X + \text{Number of tokens placed on color } Y}{\text{Number of tokens placed on color } X},$$

$$\text{Odds on color } Y = \frac{\text{Number of tokens placed on color } X + \text{Number of tokens placed on color } Y}{\text{Number of tokens placed on color } Y}.$$

If no participant has placed a token on a color up to a given period, we will say that the odds on that color are not determined yet. In that case, the odds on the color will be indicated as “ind.”

As soon as all participants have been informed of the odds, the market continues to the next period.

The example below describes how a market proceeds:

Period	Decision of the participant	Odds on color X	Odds on color Y
1	“place 1 token on color X ”	$1/1 = 1$	ind = indeterminate
2	“place 1 token on color Y ”	$2/1 = 2$	$2/1 = 2$
3	“place 1 token on color Y ”	$3/1 = 3$	$3/2 = 1.5$
4	“place 1 token on color Y ”	$4/1 = 4$	$4/3 = 1.333$
5	“place 1 token on color X ”	$5/2 = 2.5$	$5/3 = 1.666$
6	“place 1 token on color Y ”	$6/2 = 3$	$6/4 = 1.5$
7	“place 1 token on color Y ”	$7/2 = 3.5$	$7/5 = 1.4$
8	“place 1 token on color X ”	$8/3 = 2.666$	$8/5 = 1.6$

PAYMENT FOR TOKEN PLACEMENTS

In each of the 20 markets, the earnings of a participant depend on the winning color, on his or her decision, and on the **odds at the end of period 8**, which are called the **final odds**. The winning color is the one that corresponds to the color drawn by the assistant at the beginning of the market. Each market therefore has a winning color and a losing color.

For each market, a participant’s earnings are calculated as follows:

$$\text{Earnings} = \begin{cases} \text{final odds on the winning color,} & \text{if the participant has placed a token on the winning color,} \\ 0 & \text{, if the participant has placed a token on the losing color.} \end{cases}$$

In other words, participants who put their token on the winning color obtain earnings equal to the final odds on the winning color. Participants who put their token on the losing color obtain earnings of zero.

At the end of period 8, all participants observe a screen that indicates the final odds on the two colors and the winning color. This screen also informs each participant of his or her earnings for the market.

Continuing with the example presented earlier where the final odds on the color X are equal to 2.666 and the final odds on color Y are equal to 1.6:

If Y is the **winning color**,

- the earnings of the participant who placed 1 token on color X in period 1 are zero,
- the earnings of the participant who placed 1 token on color Y in period 2 are equal to 1.6 tokens,
- the earnings of the participant who placed 1 token on color Y in period 3 are equal to 1.6 tokens,
- the earnings of the participant who placed 1 token on color Y in period 4 are equal to 1.6 tokens,
- the earnings of the participant who placed 1 token on color X in period 5 are zero,
- the earnings of the participant who placed 1 token on color Y in period 6 are equal to 1.6 tokens,
- the earnings of the participant who placed 1 token on color Y in period 7 are equal to 1.6 tokens,
- the earnings of the participant who placed 1 token on color X in period 8 are zero.

If X is the **winning color**,

- the earnings of the participant who placed 1 token on color X in period 1 are equal to 2.666 tokens,
- the earnings of the participant who placed 1 token on color Y in period 2 are zero,
- the earnings of the participant who placed 1 token on color Y in period 3 are zero,
- the earnings of the participant who placed 1 token on color Y in period 4 are zero,
- the earnings of the participant who placed 1 token on color X in period 5 are equal to 2.666 tokens,
- the earnings of the participant who placed 1 token on color Y in period 6 are zero,
- the earnings of the participant who placed 1 token on color Y in period 7 are zero,
- the earnings of the participant who placed 1 token on color X in period 8 are equal to 2.666 tokens.

PRIVATE INFORMATION

To help you making your decisions you will have some partial private information about the winning color the assistant has randomly drawn at the beginning of the market. This private information will be revealed to you at the beginning of the period in which you place your token. Your private information is the result of a random draw from an urn. The urn from which your private information is drawn contains 4 balls and the exact composition of the urn depends on the winning color the assistant has drawn at random at the beginning of the market:

- If the winning color is color X , the urn from which your private information is drawn at random contains 3 balls of color X and 1 ball of color Y ,

- If the winning color is color Y , the urn from which your private information is drawn at random contains 3 balls of color Y and 1 ball of color X .

Your private information will therefore consist of either one ball of color X or one ball of color Y , but you will not know which urn it is drawn from. At the beginning of the period in which you place 1 token, you will proceed with the random draw of your private information by clicking on a button on your screen. The result of the draw will be displayed on your screen. **You will be the only one to know the result of your own private information draw.** Similarly, each of the other participants receives partial private information at the beginning of his or her period of placement.

BONUS: MAKING PREDICTIONS

Over the course of each market, in addition to your choice “place 1 token on color X ” or “place 1 token on color Y ,” you must submit a certain number of predictions. During each period, the computer will ask you for some predictions. You must indicate what is, according to you, the percentage chance that color X is the winning color as well as the percentage chance that color Y is the winning color. These two percentages make up your prediction. The percentages that you enter must be expressed in the form of whole numbers and the two percentages must add up to 100%. Recall that you will be asked to “place 1 token on color X ” or “place 1 token on color Y ” in only one predetermined period during the market, while you will be requested to make a prediction in every period. The predictions of a given participant will not be observed by other participants. The way in which you will earn money for your predictions is summarized in the table included in the appendix of these instructions. The table has been constructed on the basis of a mathematical formula that it is not necessary to know. Remember simply that your *bonus earnings* from your predictions are maximized if you honestly indicate your predictions. In other words, the farther your predictions are from what you really think, the fewer tokens you will obtain on average (the proof is available and you may ask for it at the end of the experiment). In each period, your bonus earnings will be between 0 and .15 euros.

CONVERSION OF TOKENS INTO EUROS

At the end of the experiment (after everyone has participated in the 20 markets), the assistant will randomly draw two of the markets for which your token placement will count toward your final payment. The procedure for conducting the random draw will be the following: The assistant will have 20 tickets on the table in front of him, face down. The tickets will be numbered from 1 to 20. The assistant will then randomly choose two tickets out of the 20. The two numbers will determine which two markets count toward your final payment.

Each market among the 20 conducted during the experiment has the same percentage chance to be drawn at random to count toward your final payment. The total number of tokens that you have earned in those two markets will be converted into euros at a rate of **6 euros for 1 token**. Thus, if you have obtained 5 tokens in the two markets, your final payment will be $6 \times 5 = 30$ euros. All of your predictions in each of the 20 markets will count toward your earnings.

At any time during the experiment, you may access some information about the markets you have previously participated in (the odds on each color, your placement, your private information, and your earnings) by clicking on the *history button*.

If you would like to ask a question, raise your hand and an experimenter will come to you and answer it individually. Otherwise, please do not talk during the experiment.

Before the experiment, please complete a small questionnaire on your computer. This is to verify your comprehension of the procedure. Afterwards, you will participate in a practice market that will not count toward your earnings. You will not make any of your own decisions in the practice market. Instead, the assistant will announce the decisions that you will be asked to enter during the practice market. The practice market will allow you to familiarize yourself with the environment described in these instructions.

GOOD LUCK!

APPENDIX

When the computer at which you are seated prompts you to indicate your prediction for a given period, you must enter two percentages:

- The percentage chance that, in your opinion, the winning color is color X ,
- The percentage chance that, in your opinion, the winning color is color Y .

The percentages that you enter must be expressed in the form of whole numbers (no decimals and no fractions) and the two numbers must add up to 100%. The table below shows how different pairs of percentages correspond to earnings in tokens. Because of lack of space, we do not vary the percentages by all possible increments but only by 10% increments. However, when making your predictions you can enter numbers with a level of precision of 1% (for example, (98%, 2%), (63%, 37%), (46%, 54%), or (15%, 85%)).

Predictions		Earnings in euros	
Percentage chance that the winning color is color X	Percentage chance that the winning color is color Y	If the winning color is color X	If the winning color is color Y
0%	100%	0.000	0.150
10%	90%	0.029	0.149
20%	80%	0.054	0.144
30%	70%	0.077	0.137
40%	60%	0.096	0.126
50%	50%	0.113	0.113
60%	40%	0.126	0.096
70%	30%	0.137	0.077
80%	20%	0.144	0.054
90%	10%	0.149	0.029
100%	0%	0.150	0.000

Recall that **your average earnings from making predictions are maximized if you honestly indicate your predictions**. The example below will, we hope, convince you of this.

Example: If you think that the percentage chance that the winning color is color X equals 10% and the percentage chance that the winning color is color Y equals 90%, you will earn 0.029 euros if the winning color is color X and 0.149 euros if the winning color is color Y . **The earnings you will receive on average are equal to $10\% \times 0.029 + 90\% \times 0.149$, which is approximately **0.137 euros**.**

If, in contrast to an honest prediction, you enter that the percentage chance that the winning color is color X equals 80% and the percentage chance that the winning color is color Y equals 20%, your earnings will be 0.144 euros if color X is the winning color and 0.054 euros if color Y is the winning color. Nevertheless, in reality, you believe that the percentage chance that color X is the winning color equals 10% and the percentage chance that color Y is the winning color equals 90%. Then the earnings that you will receive if your belief is correct are on average $10\% \times 0.144 + 90\% \times 0.054$, which is approximately 0.063 euros.

Therefore, by entering a prediction different from your honest one, you would lose earnings equal to 0.074 euros ($0.137 - 0.063$) for the given period. Naturally, if the predictions that you enter on your computer are even farther away from what you really believe, you would lose even more money on average.