DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

THE DESIGN AND TESTING OF INFORMATION AGGREGATION MECHANISMS: A TWO-STAGE PARIMUTUEL IAM

Kevin A. Roust
University of California, San Diego

Charles R. Plott
Califonria Institute of Technology

SOCIAL SCIENCE WORKING PAPER 1245

December 2005
Revised October 2006

The Design and Testing of Information Aggregation Mechanisms: A Two-Stage Parimutuel IAM

Kevin A. Roust   Charles R. Plott

Abstract

The research reported here is focused on the design of new information aggregation mechanisms. These are competitive processes designed for collecting and aggregating dispersed information held in the form of impression and belief, that might otherwise be impossible to get. The research explores alternative institutional forms of IAMs and how they work.

That specially designed markets can aggregate information is well documented in the literature as are the problems encountered when parimutuel-type systems are employed in an information aggregation capacity. This research is focused on new mechanisms, unlike any found evolving naturally, which mitigate these problems. These new mechanisms speed the process through which information is revealed, reduce deceptive behavior and reduce the instances of substantially incorrect aggregation (i.e., bubbles). The paper finds that a special, "two-stage" parimutuel mechanism is an improvement over previously studied parimutuel mechanisms. The two-stage parimutuel, on average, makes a prediction closer to that predicted using all available information. The mechanism suffers from fewer mirages (bubbles) than do previous parimutuel structures and it produces indicators for assessing the reliability of the information produced.

The Design and Testing of Information Aggregation Mechanisms: A Two-Stage Parimutuel IAM

Kevin A. Roust
University of California, San Diego

Charles R. Plott
California Institute of Technology

1. Introduction[1]

This research is motivated by the discovery that a properly designed information aggregation mechanism (IAM) can collect and aggregate information that is otherwise dispersed across individuals in the form of hunches and intuition. Furthermore, such mechanisms have applications.[2] The possibility of exploring the institutional range of IAMs was initiated by Plott, Wit, and Yang (2003) (PWY) who discovered that parimutual betting systems can serve such functions. By taking an institutional design approach, exploring institutions that might have no similarity with those that have evolved naturally, the research reported here attempts to improve such capabilities.

That parimutual mechanisms should have the information aggregation capacities of markets is not obvious. The structure of the betting system allows nearly unlimited opportunities for participants to conceal their information (by timing their purchases) and/or actively mislead other participants (with "bad" purchases). Indeed, as reported by PWY the betting systems can be wrong and mistakenly predict that events have high

[1] The financial support of the National Science Foundation and the Caltech Laboratory for Experimental Economics and Political Science is gratefully acknowledged. We thank Boris Axelrod and Ben Kulick for many helpful insights during the formative stages of this research.
[2] The first development and application, Chen and Plott (2002), focused on a selection of sales people engaged in forecasting sales for a Hewlett Packard product through a complete set of state dependent instruments. The general possibility of the design and use of such mechanisms was first explored by Plott

probability when they will actually occur with very low probabilities. The research reported here begins with such inaccuracies and adds institutional features to the betting system that have the potential (in theory) to avoid previously discovered problems and thereby improve the ability of the system to aggregate information.

The basic institution of interest is a “Two-Stage Parimutuel Mechanism.” While its features and the data that lead to its invention are described in detail in later sections, its features and operation have an intuitive flavor. The Two-Stage has a specially designed first stage of betting that is followed by a parimutuel betting system. The first stage is similar to a simple lottery for a fixed prize. Agents buy tickets with fixed budgets of fiat. Ticket prices are fixed in fiat and the fiat has no value other than as a means to purchase the lottery tickets. Thus, there is an incentive for individuals to spend all of their budget and in doing so release information to the system. The number of tickets purchased is not disclosed until the “ticket windows” are closed for further purchases. Thus, during the time of ticket purchases individuals learn nothing from the odds or the patterns of purchases by others. The theory here is that the bets placed by agents reflect their private information, and thus the degree to which individuals hold similar information and thus the reliability of the information that exists in the aggregate would be reflected in the pattern of bets. The second stage of the Two-Stage mechanism is a parimutuel with a clock feature that is used to adjust the price of “tickets” on the various events predicted. As time progresses ticket prices increase, so the return to ticket purchases decreases. This provides an incentive for bets to be placed early thus revealing any information that they might contain while there is still time for others to react.

The questions posed for the data are whether or not the mechanism performance satisfies the two basic principles of design methodology. The questions form the background approach to data analysis. The first question is focused on a proof of principle. Does the mechanism do what it is designed to do? As will be demonstrated in the body of the paper, the answer is that it does. (i) The Two-Stage parimutuel does aggregate

---

(2000). That properly designed markets can collect and aggregate information was discovered by Plott and Sunder (1988).

information. It does so better than previous parimutuel systems. (ii) The Two-Stage parimutuel reduces the number of mirages and changes the character of those that remain. (iii) The Two-Stage parimutuel produces indicators of when its predictions are less reliable. The second question addresses the issue of design consistency. Does the mechanism have elements of design consistency in the sense that its behavior is understandable in term of theory?. As will be outlined in later sections, its performance is understandable in terms of basic principles of economics.

The paper consists of eight sections. The first is this introductory Section One, which is followed by a brief discussion of background research in Section Two. Section Three is a detailed description of the mechanism. Section Four is a discussion of the nature of the information to be aggregated. Section Five is the information environment implemented in the testbed experiments. Section Six lists the models and measurements used to test both performance (the proof of principle) and underlying theory of how it might operate (design consistency). Section Seven contains the results and Section Eight is a summary of the conclusions.

2. Background Research and Results

The question pursued here is the basic question posed by PWY: whether or not betting systems could have properties similar to the rational expectations and information aggregation properties of markets first discovered by Plott and Sunder (1988) and generalized as information aggregation tools, see Plott (2000).[3] Nature chooses a state and individuals independently receive signals conditional on that state. The question posed is whether or not the odds that emerge from a betting process have important characteristics of the probability distribution over the possible states that results from the pooling of information held by all individuals. The distribution derived from the pooling of all observations is called the "aggregate information available," AIA. The answer given by PWY is that the odds do have important properties of AIA but frequently the odds overstate the probability of rare events and sometimes a type of “mirage” or “bubble” can occur, a case in which the odds place a high probability on a state while the AIA indicates that the probability is actually low. In other words, the information held by the individuals occasionally becomes “morphed” into something that is far from the facts while at the same time appearing to give accurate predictions, hence the word “mirage.” The two features, overstatements of the probability of rare events and mirages, are exceptions to an otherwise reliable ability to predict events even when compared to AIA.

Three features of the PWY mechanism were thought to contribute to the inaccuracies. First, there was a tendency for some individuals to avoid placing bets. Perhaps reflecting risk aversion, they preferred to keep the cash. Thus, some of the information available to subjects was not exposed to the system. Secondly, there was a tendency in PWY for bettors to place bets at the last second thus giving little or no opportunity for betting information to become aggregated in the decisions of others. The incentive for this is obvious since those betting later have the advantage of having observed the behavior of the early betters. Third, the probabilities of rare events as predicted by the mechanism were too high. PWY thought that the excessive odds on rare events were due to disequilibrium betting. In the PWY parimutuel system, a bet placed remains even though the beliefs of the bettor might change in response to the bets placed by others.

Subsequent research by Axelrod, Kulick, Plott, and Roust (2004) made changes to correct two problems they detected. First, they felt that poor performance could be traced to late betting. Secondly they felt that disequilibrium betting distorted odds, creating a “long shot bias.” Thus, AKPR first implemented incentives for agents to place bets early. This was done by imposing an increasing price of “tickets,” the claims to portions of the prize purse. That is, as time progressed the return to a bet, given the odds, was systematically lowered by increasing ticket prices, thereby creating an incentive to place bets early. Secondly the problem of disequilibrium bets was addressed by the implementation of a

[3] The design of aggregation mechanisms and application to a field setting is first introduced in 1996 by Chen and Plott (2002). That study utilized a complete set of Arrow Debreu securities as a mechanism and

dual parimutuel system. In this dual system, a parimutuel system was conducted in which bets were placed, the betting window was closed and the odds were revealed but the winning state was not revealed. Then with the information contained in private signals unchanged a second, repeat betting system was opened in which a separate prize was awarded. Bets were placed based on the same private information as before but also on the public information made available in the odds of the first exercise. Since the "equilibrium odds" were revealed in the first exercise there should be few disequilibrium trades in the second, repeat betting exercise.

The parimutuel mechanism implemented studied by AKPR produced odds that more accurately captured the AIA than did the system studied by PWY. In the repeat-betting phase, the excessive odds for rare events disappeared suggesting that the phenomenon was due to disequilibrium trades as they suspected. However, the mirages did not disappear in the second phase and indeed became even more pronounced when they appeared. Basically, when the system was bad it was very bad.

A close look at these bubbles produced by the AKPR mechanism found characteristics that lead to the design implemented and tested here. First, individual behaviors suggested that individuals sometimes sensed the existence of a mirage. During bubble periods there was a tendency for individuals to bet less. During mirage periods the variance of odds during odds development was higher than in non-bubble periods. Secondly, mirages tended to occur when the AIA was relatively poor in the sense of having a relatively large variance (a relatively flat distribution). These two facts suggested that if individuals were made aware that the AIA was poor, that fact would be incorporated into the odds development and thus remove the misleading information inherent in mirages.

These insights lead to a "Two-Stage" parimutuel information aggregation mechanism. The Two-Stage has a specially designed first stage of betting that is followed by a parimutuel betting system. The first stage is similar to a simple lottery. First, the prize is fixed and agents buy tickets with fixed budgets of fiat. Ticket prices are fixed in fiat and

---

applied it to a sales forecasting problem.

the fiat has no value other than as a means to purchase the lottery tickets. Thus, there is an incentive for individuals to spend all of their budget and in doing so release information to the system. The number of tickets purchased is not disclosed until the "ticket windows" are closed for further purchases. Thus, during the time of ticket purchases individuals learn nothing from the odds or the patterns of purchases by others. If the bets that agents place reflect their private information, then the degree to which individuals receive similar information and thus the reliability of the information that exists in the aggregate, would be reflected in the pattern of bets.

When the second parimutuel stage opens, the agents are aware that the information to be aggregated is poor, as revealed by the first stage, and they might be more likely to rely on their own information and be more attentive in interpreting the behavior of others. Such a process might be capable of engaging the capacities of participants to produce a more reliable overall aggregation of the information that exists to be aggregated.

3. The Two-Stage Parimutuel Information Aggregation Mechanism (IAM)

The mechanism is discussed in three parts. The first is a simple overview. The second describes the central parimutuel betting process used in the second stage of the mechanism. In many respects, it is the most complex and is the technical heart of the mechanism. Following that, the first stage is discussed.

The Mechanism and Information Overview

Many features of the information aggregation mechanism (IAM) and the experimental testbed operate in the research. A brief overview will be useful.

A particular experiment consisted of a number of periods. Each period contained six major phases:

1. Random selection of "true state" for the period (revealed only to experimenter) from the six possible states {A, B, C, D, E, F}.
2. Distribution of private information about the "true state" to each subject. The private information, a private signal determined independently for each individual, was a sample of three independent draws from a distribution conditional on the true state.
3. First stage of the mechanism opened for purchases of tickets and then closed. Different states were represented by different tickets. Tickets for the true state paid the holder a percentage of a fixed prize that was determined by the number of such tickets sold.
4. Total tickets sold for each state during first stage was revealed to all subjects.
5. Second stage of the mechanism opened for ticket purchases and then closed, with continuous updates as to the number of tickets sold for each state and increasing ticket prices as the period progressed.
6. "True state" is publicly announced and earnings information from both stages was distributed. Earnings were determined in the parimutuel manner described in the next section.

The Second Stage of the IAM: The Parimutuel Betting System

Plott, Wit, and Yang (2003) considered a parimutuel betting system that is based loosely on real-world markets for betting on horse races. Except as noted, the mechanism used here is the same. Each subject was provided with a quantity of experimental money ("francs," converted to U. S. dollars at the end of the experiment) that could be used to purchase "tickets" corresponding to the six possible states. The experimenter sells the tickets prior to the announcement of the winning state.

The price of the tickets increases at a known rate until the ticket sales stop at a known time. In particular, the market was open for two minutes and the ticket price would increase at a rate of one franc per second from an initial 100 francs per ticket. The time remaining before closing and the current price was reported publicly at all times.

After the ticket sales stopped, the “true state” would be announced and earnings would be based on individual ticket purchases and the "true state" in the following manner.

First, the winnings for each "true state" ticket would be calculated:

$$\text{Per Ticket Winnings} = \frac{\text{Total Revenue from Ticket Sales} + \text{House Bonus}}{\text{Total Number of Winning Tickets}}$$

Then an individual earns:

$$\text{Number of Winning Tickets Held} * \text{Per Ticket Winnings} + \text{Unspent Experimental Money} - \text{Loan}$$

As a period progressed, the total tickets sold and "Per Ticket Winnings" (conditional on the market ending instantly and the given state winning) were reported publicly for each state.

The First Stage of the IAM:

The first stage of the IAM was designed to make public whether the information that exists in the system was decisive or ambiguous. The mechanism is based on the hypothesis that bets placed in the absence of knowledge of how others bet and with incentives to place bets rather than withhold and placed without any coordinating device through which attempted deceptions could be coordinated would reveal the degree to which a consensus exists about the true state. The structure reflected that theory.

The betting windows opened for a fixed period during which subjects bought tickets associated with the various states at fixed prices and without knowing the amounts purchased by other subjects. Winning tickets gave the holder a percentage of a fixed prize. Specifically, let T_{iw} be the tickets held by individual i for winning state w and let P represent the size of the prize associated with the winning state. Then the earnings of individual i are: $(T_{iw}/ \Sigma_j T_{jw})$ P.

Purchases were made with fiat that had no value outside the purchases of tickets for the period at hand. Each agent had a fixed budget of fiat. This effectively increased the cost of keeping silent during the first stage and was expected to induce more participation. Further, this had the potential to reduce subjects' informational size – as studied by Roust (2004), informational size in this market is inversely related to the number of participating subjects. If ten subjects are participating in the market, then any individual subject controls about 10% of the ticket purchases; if five subjects drop out then the remaining subjects each control about 20% of the ticket purchases and thus have more influence over the aggregate distribution of ticket sales.

When tickets sales stopped the total of all ticket sales and the percentage of sales for each state were made public.

After ticket sales were made public, the ticket sales began for the parimutuel, second stage of the mechanism. After the sales of tickets closed for the second stage, the true state was announced and subjects were informed of their earnings in both stages.

4. The Information Environment[4]

Each period started with a computer selecting one of the six possible states from a uniform distribution (with 1/6 probability for each state). This selection is indicated to the experimenter and is used as the "true state" for that period.

The period continued with the computer generating three independent signals of the state for each subject. That is, each subject was provided with three independent signals as to this state and then the agents were allowed to interact in a particular market setting. Their purchases in the market and the true state of the world then determined their earnings.

These signals correspond directly to the possible states of the world and are drawn from {A, B, C, D, E, F}. The signals were generated with a 1/3 probability of receiving a correct signal (signal matches true state) and a 2/15 probability for each incorrect signal (signal is different from the true state). If the true state were A, then it would be as if the signals were drawn independently (with replacement) from an urn containing fifteen balls labeled as follows:

A A A A A B B C C D D E E F F

For analysis purposes, it is occasionally useful to separate the agents into three groups based on the apparent quality of their private information. A subject receiving three matching signals (AAA, BBB, CCC, DDD, EEE, or FFF) has the strongest private information. A subject receiving three different signals (e.g., ABC) has received the weakest private information. A subject receiving exactly two matching signals (e.g., AAB) lies between those extremes.

Since each subject receives three independent draws, their private information has a multinomial distribution with six possible outcomes and three trials using the probabilities described above. This distribution has the following joint probability mass function:

$$f\left(x_A,\ldots,x_F \middle| \text{True State} = S\right) = \frac{m!}{x_A!\cdots x_F!}\, p_C^{x_S}\, p_N^{m-x_S}$$

$$x_i = \text{Number of signals matching i}$$

$$m = \sum_{i=A}^{F} x_i = \text{Total number of signals}$$

$$m = 3 \quad p_C = \tfrac{1}{3} \quad p_N = \tfrac{2}{15}$$

Bayes' Law can then be used in the standard way to identify the posterior probability distribution over the states given the private signal.

[4] This is the "Probabilistic Information Condition" (PIC) structure used by Plott, Wit, and Yang (2003).

$$P(S|x) = \frac{f(x|S)P(S)}{\sum_{i=A}^{F} f(x|i)P(i)} = \frac{f(x|S)}{\sum_{i=A}^{F} f(x|i)}$$

(since $P(i) = \frac{1}{6}$ for all states)

These posterior probabilities can be easily summarized in a simple table like Table 1.

Table 1. Posterior probability distribution across states given private information

	Most Frequent (e.g., A)	2nd Most Frequent (e.g., B)	3rd Most Frequent (e.g., C)	Each Not Drawn (e.g., DEF)
Three matching draws (e.g., AAA)	75.8%	4.8%	4.8%	4.8%
Two matching draws (e.g., AAB)	49.0%	19.6%	7.8%	7.8%
Three different draws (e.g., ABC)	23.8%	23.8%	23.8%	9.5%

A table like this was provided to subjects (and explained) in the instructions for the experiment, to build confidence that the subjects developed reasonable beliefs based on their private information.

5. Experiment Design and Procedures

Five experimental sessions were conducted using the Two-Stage parimutuel structure. Most of the parameters were held constant across all experiments and all rounds, with only the true states and private information varying.

There was only one major difference among experiments. The first two experiments (020305 and 020312) allowed subjects to keep any unspent francs from the first stage (and be paid for them at the end of the experiment). In the later three experiments, subjects were not paid for any unspent francs from the first stage (although they were still paid for unspent francs from the second stage).

Experiment parameters are listed in Table 2.

Table 2. Experiments and Experimental Conditions

Experiment	020305	020312	020329	020423	020507
Subjects	10	15	10	10	10
Practice Round #s	1	0	0	0	0
Non-Practice Round #s	2-11	1-20	1-20	1-20	1-20
Stage A Length	60 sec.	60 sec.	60 sec.	60 sec.	60 sec.
Stage A Ticket Price	100	100	100	100	100
Stage A Budget	10,000	10,000	10,000	10,000	10,000
Stage A Loan	5,000	5,000	0	0	0
Stage A Bonus	20,000	20,000	20,000	20,000	20,000
Stage A Earnings	Winnings + Balance – Loan	Winnings + Balance – Loan	Winnings Only	Winnings Only	Winnings Only
Stage B Length	120 sec.	120 sec.	120 sec.	120 sec.	120 sec.
Stage B Ticket Prices	100-220 (+1f/sec)	100-220 (+1f/sec)	100-220 (+1f/sec)	100-220 (+1f/sec)	100-220 (+1f/sec)
Stage B Budget	50,000	50,000	50,000	50,000	50,000
Stage B Loan	20,000	20,000	25,000	25,000	25,000
Stage B Bonus	100,000	100,000	100,000	100,000	100,000
Stage B Earnings	Winnings + Balance – Loan	Winnings + Balance – Loan	Winnings + Balance – Loan	Winnings + Balance – Loan	Winnings + Balance – Loan
Exchange Rate (francs : cents)	200:1[5]	250:1	250:1	250:1	250:1

Subjects and Recruiting

The subjects were Caltech undergraduates, and they were generally currently enrolled in economics courses. Although they generally did not have experience with either this

[5] The exchange rate was announced to be 400 francs : 1 cent at the beginning of the experiment (in secret anticipation of completing 20 rounds). After the experiment ended, it was announced that the rate was increased to 200 francs : 1 cent.

particular experiment or other parimutuel experiments, they did have previous experience as subjects in the laboratory.[6]

Instructions and Training

Subjects were given written instructions (which are available in the appendix). These instructions were read to the subjects with the aid of presentations on a whiteboard – items such as a drawing of an urn with states in it, the posterior probability table included above, and the earnings calculation formulae. After the instructions were read and questions were answered, a practice round was opened for the subjects. This round was identical to later rounds, except that no real money was involved. This round allowed subjects to familiarize themselves with the software used and the calculations required for the remaining rounds.

Technology

A purpose-built piece of software was used for all communication and transactions during the experiment. This software is the same as that used for the Axelrod, Kulick, Plott, and Roust experiments, save for minor modifications to allow for the two-round structure (without publicly reported sales during Stage One).[7] The software was used to generate the true state and private information (on the experimenter's interface); distribute the private information to subjects (through each of their screens); receive orders from subjects while the market is open; report (with each transaction by any subject) the current time in the round, the current ticket price, personal francs remaining, personal holdings of tickets, aggregate holdings of tickets (by state), and current per ticket winnings (by state, as if that state were winning); and provide a report of current and total earnings after each round. This software was extensively tested before any experiments were performed.

[6] The experiments reported here all occurred in the spring of 2002. Experiment 020305 and 020507 had one subject in common. Experiment 020305 also had two subjects with experience from the Axelrod, Kulick, Plott, and Roust experiments (in 2001). Experiment 020312 had five subjects with AKPR experience. Presumably none of the subjects participated in the Plott, Wit, and Yang experiments (from 1995), although this has not been verified.

6. Information Aggregation Measurement And Models

The information revealed by the IAM is typically assumed to be represented by the observed odds. The "Observed Odds" of a state refers to the percentage of all tickets sold that correspond to that state. That is, the observed odds of a state s are in the ratio (Total Spending/Tickets for s). In a sense, the inverse of the odds can be interpreted as a probability. To connect this measure with behavior assumes that subjects act to maximize expected return; they will purchase tickets so that $\text{Prob}(s)\times\frac{\text{Total Spending}}{\text{Tickets for s}}$ is constant across all states where they purchase tickets. In particular, this implies that the distribution of tickets reflects the probabilities predicted in a market of expected-return-maximizing subjects. These are the "implicit prices" defined by Plott, Wit, and Yang (2003). In this sense, the odds reflect the "group probability" distribution across the states and the questions to be addressed are how "close" is this probability to the aggregate information available and how is it related to individual behavior.

To measure how close are the odds to other distributions, such as the aggregate information available, we use a measure described in Dacunha-Castelle (1977, 49-50) and rediscovered in Würtz (1997).[8] The concept is that of a "displacement function" of two distributions *X* and *Y* defined as half of the L^1-norm difference between them. In particular, the ***half-L^1 distance*** is half of the absolute deviation between X and Y:

$$D(X,Y)=\tfrac{1}{2}\sum_{s}\left|X_s-Y_s\right|.$$

The half-L^1 distance measures the distance between two distributions as the difference in probability mass between them. If the probability distribution reflected in the odds is *X* and if some other distribution, say AIA, is *Y* then D(Odds,AIA) is a measure of how close the distributions are to one another. Dacunha-Castelle and, independently, Würtz prove

---

[7] The software was originally written by Boris Axelrod. The modifications were primarily by Hsing-Yang Lee, with additional minor modifications by Kevin Roust.
[8] The reference to Dacunha-Castelle was called to our attention by Würtz .

that the half-L^1 distance is an upper bound on the error caused by using the distribution Y for statistical analysis when the true distribution is X.[9]

Typically, the analysis will inquire whether or not information has become aggregated by the mechanism and is thus reflected in the odds. Such analysis requires a baseline of what would be the case if no information has been aggregated. The baseline models can also serve as models of behavior and will be used as such.

Equally Likely (EL) Model

The EL model assumes that the posterior probabilities are exactly the prior probabilities (of 1/6 per state). If applied as a behavioral model, it is as if a subject using this model will purchase equal numbers of tickets for each state. If all subjects follow the EL model, then the aggregate ticket sales should be equal among states as well. It is as if subjects ignore all market information and ignore all of their private information as well.

Decision Theory (DT) Model

The DT model uses individual private information to identify the most likely state. If a subject received two or three matching signals, then the most likely state corresponds to the matching signals. If they received three different signals, then the three states represented are equally likely and more likely than the other three states. Having identified the most likely state(s), these subjects will then make purchases only in that state. The aggregate prediction of the DT model is then the average across subjects of these purchases. This model assumes that subjects ignore all market information and ignore the Bayesian posterior distributions as well (although they do use the raw signals).

Average Beliefs (AVB) Model

The AVB model uses the Bayesian posterior distribution derived from individual private information. As a baseline it takes the average of these as a measure of no aggregation in

[9] For example, if someone wanted to make a prediction using all available information, they might perform a statistical test on the AIA distribution with significance level p. Equivalently, they could use the Observed Odds and perform the test with significance (p - D(AIA,Odds)) and achieve the goal of having a significance p test under AIA.

the market. As a behavioral model it assumes that subjects will make purchases with this distribution. For example, a subject who receives three matching signals will spend roughly 3/4 of their available francs on the corresponding state and roughly 5% on each of the other states, as in Table 1. As with the DT model, these purchases are then averaged across subjects to generate the prediction of the AVB model. This model assumes that subjects ignore all market information but do use their private information to best effect.

Aggregate Information Available (AIA) Model

The AIA model represents maximal information aggregation. The AIA model uses all private information (three signals for each subject), pools the data, and then uses Bayes' Law to predict the state.

Equilibrium Model

A full characterization of the equilibria of the parimutuel market still escapes the authors. However, a few equilibria have been identified in Roust (2004). These are characterized by a set of sufficient conditions for the previous models to be the equilibrium aggregate ticket sales.

The AIA model can be supported as a Nash equilibrium in a single-stage parimutuel (before the true state is revealed) if risk neutral subjects develop common beliefs ρ. If the subjects purchase tickets in proportion to their common beliefs, then per-ticket winnings will be proportional to $1/\rho$, meaning that the expected return to a ticket purchase is identical in every state. With identical expected returns and risk neutrality, subjects are indifferent among ticket allocations, permitting purchases in proportion to common beliefs (which could be the AIA distribution).

The DT model can be supported as a Nash equilibrium in a single-stage parimutuel (before aggregate ticket sales are revealed) if Bayesian risk-neutral subjects ignore the effects that ticket sales (both theirs and others') have on per ticket winnings.[10] In this case, subjects will believe that the per ticket winnings will be identical for all states, meaning that expected returns will be maximized by investing only in the most likely state given their information. This corresponds with the DT model.

Entropy

We use two measures of the "quality" of the aggregate information. The first is the ***entropy*** of the AIA distribution. Entropy is a measure of the "flatness" of a distribution. In particular, entropy of a probability distribution is defined as $-\sum_s \text{Prob}(s) \cdot \ln(\text{Prob}(s))$.

A distribution with an entropy of 0 places probability 1 on one particular state. A distribution with an entropy of ln(6) (about 1.79) places probability 1/6 on each of the six possible states. When comparing distributions, higher entropy numbers are associated with flatter distributions. "High" and "low" entropy refer to "more flat" and "more peaked" distributions, respectively. Entropy of the AIA model can be used to measure the "quality" of aggregate information.

AIA Mode Size (Most Likely Probability)

The second measure is and the probability that the AIA model assigns to the most likely state. The ***AIA mode size (the most likely probability)*** is a measure that indicates how strongly the private signals agree, and theoretically varies from 16.7% (all states equally likely) to 100.0% (one state is certain).[11] Due to the structured nature of these Bayesian probabilities, the probability-based and entropy-based measures have similar ordering properties (the Spearman rank correlation is 0.991). In particular, since the analysis that follows simply divides the rounds into bins, based on the chosen measure, there will be little difference in the categorization between an entropy-based ranking and a probability-based ranking.

---

[10] Much like in the standard competitive model.

[11] The observed range is 26.1% in 020329 Round 6 to 99.9996% in 020312 Round 3. 020329 Round 6 had 10 subjects who collectively received 6 As, 6 Ds, 6 Fs, 5 Bs, 5 Cs, and 2 Es for a tri-modal AIA prediction (26.1% A, D, F; 10.5% B, C; 0.7% E). 020312 Round 3 had 15 subjects who collectively received 21 Cs, 7 As, 6 Es, 4 Bs, 4 Fs, and 3 Ds for a very peaked AIA prediction (99.9996% C; 0.0003% A; 0.0001% E; 0.00002% B; 0.00002% F; 0.000007% D). One of the AKPR experiments (010510 Round 3) reached 99.9999% with 26 Fs out of 57 total signals (19 subjects).

Weak and Strong Mirages: Market Mistakes

Two sorts of market mistakes are identified in this paper. The first of these mistakes is called a *weak mirage* and is a period where the most likely state according to the Observed Odds is not one of the most likely states according to the AIA model. In particular, $\max_{s} Odds(s) > \max_{\substack{\text{s maximal}\\ \text{in AIA dist}}} Odds(s)$. A market suffering from a weak mirage is failing to provide even the most basic information that one would hope to receive from an information aggregation market – it is not identifying the most likely state.

The second, more serious, mistake is called a *strong mirage* and is a period where the market gives the most likely state according to the Observed Odds at least twice as much weight as the most likely state according to the AIA model. In particular, a strong mirage occurs when a weak mirage occurs and $\max_{s} Odds(s) > 2 * \min_{\substack{\text{s maximal}\\ \text{in AIA dist}}} Odds(s)$.[12] "Twice" is an arbitrary choice, but does provide some separation between slightly confused markets (that barely choose the "wrong" state while putting significant weight on the "right" state) and very confused markets (that put most of their weight on the "wrong" state at the expense of the "right" state). A market suffering from a strong mirage fares even worse that one suffering from a simple weak mirage – not only did the market fail to identify the most likely state, but it gave that state a very low likelihood of occurring.

7. Results

The methodological approach for analyzing data from testbed experiments turns on two fundamental questions (Plott, 1994) that are addressed in the two subsections of the section. First, does the mechanism do what it is supposed to do? This question asks for a proof of principle (or proof of concept), and is addressed in the first of the subsections.

[12] If the AIA-model is unimodal, this definition is easy to interpret: a strong mirage occurs when the Odds mode has Odds that are at least twice the Odds of the AIA mode. When AIA has multiple modes, however, a strong mirage occurs when the Odds mode has Odds that are at least twice the Odds of one AIA mode. An alternate definition ("twice the Odds of all AIA modes", using "max" rather than "min") would weaken one of the thirteen Stage Two strong mirages (020507 round 3).

Second, does it work for understandable reasons? This question asks about design consistency in the sense that the mechanism works according to the principles that led to the design in the first place. Unless one has some assurance that any success exhibited by the mechanism is guided by robust principles, one has no basis for expecting that the mechanism will continue to work satisfactorily when implemented in a field environment in which the parameters can be substantially different from those present in the testbed. It is design consistency that gives confidence that the move from the simple laboratory environment to a complex field environment will be successful. This second question is addressed in the second subsection.

Proof of Principle

The first four results say that the mechanism works to aggregate information and it does so better than previously examined mechanisms. Information aggregation is taking place (Result 1). The information aggregation is better in the Two-Stage Parimutuel mechanism than in other parimutuel mechanisms that have been tested (Result 2) and the mechanism produces fewer mirages than do other mechanisms (Result 3). Following those results are additional conjectures suggesting that the mechanism produces indicators of when the results might be reliable and when they are not. While statistical support for these properties does exist they are listed as "conjectures" because we think that they characterize general tendencies of the class of mechanisms and have potential for being useful in more complex applications of IAMs.

RESULT 1. Information aggregation is occurring in the Two-Stage parimutuel IAM.

Support. As indicated in Table 3, the observed odds of the two-stage, increasing price parimutuel are closer to AIA than is the prediction of any of the non-aggregation models (AVB, DT, EL). The observed odds contain more information than is available to the agents acting independently (on average). A paired t-test rejects the hypothesis that the observed odds are further from AIA than the AVB, DT, and EL models are ($p << 1\%$ in all cases). On average, the Observed Odds for Stage Two are closer to the "true" AIA odds than individual information allows.

Table 3.

Stage Two Observed Odds are Closer to the AIA Model Than Any Other Model

	half-L^1 distance from best available prediction (AIA)
Aggregate Information Available (AIA) model	0.000
Observed Odds for Stage Two	0.332
Decision Theoretic (DT) model	0.477
Observed Odds for Stage One	0.558
Average (Bayesian) Beliefs (AVB) model	0.607
Equally Likely (EL) model	0.723

Recall that a mirage is a failure of the parimutuel to identify the most-likely state. "Weak mirages" are cases where the most likely state under the AIA model is not predicted to be the most likely state by the Observed Odds. "Strong mirages" are cases where the Observed Odds are so badly in error that some wrong state is believed to be at least twice as likely as the AIA mode state is.

RESULT 2. Two-Stage parimutuel IAMs aggregate information better than previous parimutuel IAM designs do.

- **(i) Overall the Two-Stage parimutuel is more accurate than other parimutuel IAMs.**
- **(ii) Considering only the periods of non-mirage, the Two-Stage parimutuel is more accurate than the others.**

Support. For support of (i) consider all periods, including mirage and non-mirage periods. The observed odds in Stage Two of the Two-Stage parimutuel IAM are closer to the AIA model than are the Observed Odds for the PWY PIC-4, PWY PIC-2,[13] AKPR Non-Repeat,[14] AKPR Repeat,[15] and Two-Stage Stage One.[16] Table 4 contains the statistical tests. The half-L^1 distance (half of the absolute deviation) between the odds and AIA are least for the Two-Stage Parimutuel Stage Two. This means that the Two-Stage parimutuel is a better substitute for AIA than the other parimutuel mechanisms. If AIA is unknown (as it typically is in applications), then a statistical test using the Two-Stage parimutuel Observed Odds will be more accurate than the same test using the odds from another parimutuel.

Table 4. Stage Two Observed Odds are Closer To the AIA Model Than Are Other Parimutuels

Treatment	D(Odds-AIA)	Standard Deviation	Probability better than RP Stage Two
Two-Stage Stage 2	0.332	0.225	---
AKPR Repeat	0.380	0.264	15%
AKPR Non-Repeat	0.409	0.194	2%
PWY PIC-2	0.452	0.172	<<1%
PWY PIC-4	0.485	0.192	<<1%
Two-Stage Stage 1	0.558	0.137	<<1%

For support of (ii) consider periods where the parimutuel correctly predicts the most likely state (non-mirage periods). For these periods the observed odds in Stage Two of the Two-Stage parimutuel are closer to the AIA model than are the observed odds for either the AKPR non-repeat, the AKPR repeat and the Two-Stage Stage One. See Table 5.

[13] Plott, Wit, and Yang conducted experiments with a four- or two-minute lag before the session could randomly end.
[14] These are the periods reported in AKPR in which the information was not a repetition of the previous period.
[15] These are the periods in AKPR in which the information was identical to that of the previous period.
[16] These are the first part of the Two-Stage Parimutuel Mechanism.

The half-L^1 distances between the odds and AIA for non-mirage periods are least for the Two-Stage parimutuel Stage Two, and the improvement over non-repeated periods is statistically significant. This suggests that the Two-Stage parimutuel reduces disequilibrium trades at least as well as the repeated parimutuel even when both IAMs are predicting the correct state.

Table 5. Controlling for mirages, the Two-Stage parimutuel is better at aggregating information.

Treatment	D(Odds-AIA)	Standard Deviation	Probability better than RP Stage Two
Two-Stage Stage 2	0.236	0.148	---
AKPR Repeat	0.247	0.135	36%
AKPR Non-Repeat	0.317	0.135	<<1%
Two-Stage Stage 1	0.539	0.122	<<1%

RESULT 3. Two-Stage parimutuel IAMs suffer from fewer Strong Mirages than AKPR parimutuels.

Support. Weak Mirages are about as common in Stage Two of a Two-Stage parimutuel IAM as in the non-repeated and repeated parimutuel IAMs studied by Axelrod, Kulick, Plott, and Roust (2004). Strong Mirages are far less common as indicated in Table 6. While the weak mirage differences are insignificant, a Fisher-Irwin test[17] reports a p-value of less than 0.2 that the underlying probability of a strong mirage is identical.

[17] The Fisher-Irwin test tests the hypothesis that two binomial populations have the same underlying probability parameter. It is an exact test, using the hypergeometric distribution. See, for example, Ross (1987, pp. 232-233).

Table 6. Observed Frequency of Mirages under various treatments

	Weak Mirage Frequency	Strong Mirage Frequency
Two-Stage Stage 2	26%	14%
AKPR Not Repeat	31%	23%
AKPR Repeat Round	26%	26%

This result suggests that the Two-Stage parimutuel IAM has little added potency against weak mirages, but does reduce the frequency of strong mirages. This property suggests that the first stage reveals when the aggregate information is somewhat poor and thus prevents the second stage from overreacting to the aggregate ticket sales data.

The following conjectures are weakly supported in the data. Parts of these conjectures could be stated as results with weak statistical support. We choose instead to list them as a series of conjectures in order to illustrate the general tendencies that the data seem to reveal.

CONJECTURE 1. Mirages are more likely with "bad" information in Two-Stage parimutuel IAMs.

Support. As the quality of aggregate information decreases (represented by a decrease in the probability associated with the most likely event under AIA), the frequency that a mirage is observed tends to increase, as indicated in Figure 1. The probability of observing a mirage increases from 5% (with good aggregate information) to 50% (with bad aggregate information).

CONJECTURE 2. Strong Mirages in Two-Stage parimutuel IAMs are associated with "bad" aggregate information.

Support. As indicated in Figure 1, three of the thirteen strong mirages during Stage Two occurred when the aggregate information was in the better half of all observed information, as measured by AIA mode size. On the other hand, five of those strong mirages were in the worst bin. Due to small sample sizes this is not a significant effect, however.

[FIGURE 1 ABOUT HERE]

OBSERVATION. The second stage of a Two-Stage parimutuel IAM can correct for early mistakes.

Generally, the Stage One and Stage Two markets choose the same state as "most likely". In some cases, however, the two markets disagree, with the Stage Two market being more likely to be correct. Of the 90 Two-Stage parimutuel periods, the two stages disagree on the most likely state nine times – the second stage is correct four times while the first stage is correct only once. Both stages are correct in 63 other periods.

RESULT 4. Two-Stage parimutuel IAMs suffer fewer mirages than previously studied designs.

Support. Stage Two of a Two-Stage parimutuel IAM generally has fewer Strong Mirages than AKPR parimutuel IAMs. The parimutuel mechanisms behave similarly when faced with very good aggregate information, but when faced with less perfect information the Two-Stage parimutuel clearly improves on the AKPR markets (as measured by strong mirages). The rate of strong mirages is roughly 10% lower in the Two-Stage parimutuel markets, as indicated in Figure 2.

Plott, Wit, and Yang (2003) discuss a model they call "herding," which corresponds to the "weak mirage" concept used here. For their "NOT SETS" condition, they indicate that 9 of 38 periods suffered from weak mirages (PWY 2003, 344).[18] This 75% success

[18] Their discussion suggests that these herding events would also satisfy the definition of a strong mirage.

rate (of avoiding mirages) compares poorly with the Two-Stage parimutuel IAM, particularly given that the NOT SETS aggregate information assigns 100% probability to a particular state.

[FIGURE 2 ABOUT HERE]

The fact that strong mirages happen less often under the two stage mechanism and the fact the there seems to be a relationship between mirages and the quality of aggregate information as measured by AIA mode size, that exists in the environment suggests that individuals sense when the system might be in a bubble and adjust their behavior accordingly. That means that an observer that has none of the private information held by participants nevertheless might also be able to detect when the system is producing a mirage and thus when the predictions of the system are unreliable. The following results chronicle such possibilities.

CONJECTURE 3. Aggregate spending is reduced during Weak Mirage rounds, in comparison with non-mirage rounds and Strong Mirage rounds.

Support. Subjects spend more in total during non-mirage rounds than during Weak (not Strong) Mirage rounds, suggesting that participants are aware of the mirage (75% and 70% of aggregate budget, respectively). Spending during Strong Mirage rounds returns to the non-mirage level, suggesting that participants are unaware of strong mirages.

CONJECTURE 4. Purchases occur later in mirage rounds.

Support. The timing of ticket purchases is different between non-mirage and mirage rounds, suggesting that participants may be aware of the mirage. The difference makes itself apparent before one second has elapsed in Stage Two, as indicated in Figure 3 below. With respect to Weak (not Strong) Mirages, the difference starts at approximately 20% of average spending in the first second. With respect to Strong Mirages, the difference begins at a similar level, but completely disappears by the end of the period.

While this offers some hope that outside observers could use aggregate behavior to identify mirages, this difference is too small to be statistically significant. The standard deviation on aggregate spending in the first second of the second stage is roughly 19% in both non-mirage and mirage rounds – overwhelming the 6% difference. However, it is possible that this variance can be decreased through a closer examination of individual agents or with a longer series of experiments.

[FIGURE 3 ABOUT HERE]

CONJECTURE 5. In the early parts of a mirage round, the observed odds demonstrate more uncertainty than in a non-mirage round.

Support. Rounds that are not mirages show (on average) lower entropy in their observed odds than weak and strong mirage rounds. This difference is apparent after the first second, and the gap remains roughly stable throughout the round, as indicated by Figure 4. Thus, when facing a mirage, the observed odds have a "flatter" distribution than when not facing a mirage – the ticket purchases have been more evenly spread out across the states.

[FIGURE 4 ABOUT HERE]

Together, Conjectures 3, 4, and 5 suggest that an outside observer who can monitor total aggregate spending may be able to detect weak mirages (but not strong mirages). An outside observer may be able to identify both strong and weak mirages with information on the final distribution of tickets, changes in that distribution through the period, or changes in aggregate spending through the period.

Design Consistency

The results above suggest that the Two Stage Parimutuel IAM does in fact successfully aggregate information and does so better than previously developed parimutuel-based IAMs. The question now posed is whether we can understand why. In Result 5, the potential behavioral models, AIA, AVB, DT, and EL models will be compared with the Observed Odds from both Stage One (no aggregation expected) and Stage Two (aggregation expected). Each of these models is based on particular assumptions about individual agent behavior. If these models help establish system behavior in relation to the behavior of the individuals in the system we have reason to believe that the performance of the system is not due to some particular choice of parameters or procedures in the testbed and thus have reason to hope that the success of the Two-Stage parimutuel mechanism will be exhibited in different and perhaps more complex environments.

RESULT 5. The (individual Bayesian) AVB model is the best predictor of Observed Odds in Stage One of the Two-Stage parimutuel IAM, while the (aggregate Bayesian) AIA model is the best predictor for Stage Two.

Support. The AVB model most closely predicts the Observed Odds in Stage One of a round, followed by the DT, EL, and AIA models, as indicated in Table 7. Paired t-tests show that the AVB model is a significant improvement over all three other models (p << 1%). This suggests that no information aggregation is occurring during the first stage of the Two-Stage parimutuel, as expected.

The AIA model most closely predicts Observed Odds in Stage Two of a round (but its predictions are not statistically significantly closer than the predictions of the DT model). The AIA and DT models both improve significantly over the AVB and EL models (p << 1%). Further, AIA is the only model that predicts Stage Two better than Stage One. The AIA model is a significantly better predictor of Observed Odds in Stage Two than in

Stage One, while the DT, AVB, and EL models are significantly better predictors of Stage One.

Table 7. Average half-L^1 Distance between Observed Odds and various Models

	Distance from Stage One Odds	Distance from Stage Two Odds
AIA model	0.558	**0.332**
DT model	0.161	0.349
AVB model	**0.136**	0.455
EL model	0.231	0.566

Result 5 demonstrates that substantially different individual behavior occurs in the two different stages. The rational expectations model (AIA) shifts from being the worst predictor of odds in Stage One to being the most accurate predictor of odds in Stage Two. This dramatic shift in behavior is itself an indication that information aggregation is facilitated by the mechanism. Thus, the Stage Two data are understandable.

The Stage One data present a bit of a paradox. Under risk neutrality, the most accurate model should be the DT model and while the DT model competes well with AVB, it is not as accurate as AVB. One could conjecture that risk aversion plays a role in giving AVB such predictive power, but at the moment it remains unexplained.

An important consideration in this analysis is the theory of “information size” developed by McLean and Postlewaite (2001-2004), which explains why deceptive behavior suggested by game theory does not destroy information aggregation during the second stage of the process. According to the theory when agents are “informationally small” they do not engage in information based strategic behavior. The possibility that their actions reveal strategically important information to others is not reflected in the decision process of informationally small agents but such possibilities will be reflected in the decisions of agents that are informationally large. The next result says that when agents are informationally small in the McLean and Postlewaite sense[19] they do not tend to act in an informationally deceptive manner and when informationally deceptive behavior does take place it tends to be undertaken by agents that are informationally large.

By design, agents in the testbed are almost always informationally small. Thus, the success of the mechanism is understandable in the light of game theoretic reasoning that suggests that that the performance should not be successful. The result also signals a limitation of the success of the mechanism that exists in principle. The theory that explains that success of the mechanism in informationally small environments suggests that the success will not extend itself to the case in which agents are informationally large. If the system is characterized by information monopolists, oligopolists, or exists in a highly concentrated form, the mechanism might not work.

RESULT 6. Individual behavior in the Two-Stage parimutuel IAM is generally consistent with McLean and Postlewaite's Informational Size theory.[20]

Support. Roust (2004) codes each subject's behavior in Stage Two of each round as "deceptive" or not, and then tests a series of hypotheses derived from informational size theory. The basic prediction of the identified equilibrium is supported in all those regressions, including the simple logit regression reported in Table 8 (Roust 2004, 30). The table reports both the theoretically predicted probability of deceptive behavior for each of the types of private information an individual can receive and the probability of deception as measured from the data. The qualitative predictions of the theory are clearly evident in the data. When the theory predicts that the individual will not act deceptively there is very little tendency to do so and when the theory predicts deceptive behavior there is about a 25% probability that it will be observed.

[19] Agents who receive strong information, a sample consisting of three draws of the same state, is informationally large when participating in an environment with fourteen or fewer agents. All agents in all other tested environments are informationally small according to the McLean and Postlewaite definition. An agent with two matching signals would be informationally large if there were eight or fewer agents.
[20] This is reported in more detail in Roust (2004).

Table 8. Logit Regression Of "Deceptive" Behavior Dummy On Type Of Private Information. Better Information Is Associated With More Deception, as predicted by Informational Size theory.

Independent Variable	Estimated Coefficient	Standard Error	t-Statistic	Measured Prob. of Deception	Predicted Prob. of Deception
Constant ("ABC")	-3.45	0.27	-12.72	3.1%	0%
"AAB"	0.95	0.33	2.91	7.6%	0%
"AAA"	1.42	0.51	2.77	25.3%	69%
Log-likelihood		-194.36			
Likelihood Ratio		12.18 (0.5%)			

8. Summary of Conclusions

By weaving together the results the following summary can be drawn. On average the Two-Stage parimutuel information aggregation mechanism aggregates information in a manner similar to the "rational expectations" found in markets. Furthermore, the Two-Stage parimutuel mechanism aggregates information better than other parimutuel mechanisms. The improved performance over other mechanisms can be traced to the two special institutional features. First, the increasing price feature speeds the infusion of the information into the market and avoids the phenomena of last second betting. Secondly, the first stage of the process reduces the instances of mirages by warning participants when they are in an environment in which the information to be aggregated is of low quality. Participants, realizing that the underlying information quality is low, place less weight on the behavior of others when developing their beliefs. Basically, the first stage indicates to participants when the information in the environment is not strong and, as a result, the strongly misleading mirages are avoided.

While weak mirages still exist, the behavior of the system produces signals that indicate that the information produced by the parimutuel might not be reliable. These signals are contained in the pace of the betting, the variability (entropy) of the odds during the odds development phases, and the total volume of betting. When there are no mirages the fit of the prediction is better than other IAMs studied. When considered overall including mirages the fit of the prediction is still better than other IAMs studied and there are fewer mirages.

The pattern of individual behavior is understandable, suggesting that the results are constructed on a reliable set of underlying principles. When placing bets during the first stage the betting is similar to the choice of a portfolio under uncertainty. Neither price nor information about the behavior of others is available in the first stage and coordination of strategies is not possible so the individuals' actions are consistent with the theory of decisions, as should be expected. During the second stage of the mechanism, the absence of information destroying strategic behavior is explained by the fact that individuals are informationally small, since theory suggests that when that is the case decisions that reveal beliefs are made. When placing bets during the second stage, individuals behave substantially as suggested by the model of rational expectations. The fact that theory tells us that strategic behavior can destroy the information aggregation success when information size is not small signals possibly important limitations of the mechanism in applications. Of course, the fact that the theory itself is incomplete signals a need for caution in applications. The detail and the dynamics of rational expectations formation remains a mystery, just as the phenomenon remains a mystery in the context of markets.

References

Axelrod, B., B. Kulick, C. Plott, and K. Roust; 2004; "Information Aggregation in an Experimental Study: Repeated Parimutuel Markets"; working paper (dated 08 Sep 2004).

Chen, K. Y., and C. Plott; 2002; “Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem”; Social Science Working Paper 1131, California Institute of Technology.

Dacunha-Castelle, D.; 1978; "Vitesse de Convergence pour certains Problèmes statistiques"; in *Ecole d’Eté de Probabilités de Saint-Flour VII-1977 (Lecture Notes in Mathematics 678)*; ed. P. L. Hennequin; Springer-Verlag, Berlin.

McLean, R., and A. Postlewaite; 2001; "Efficient Auction Mechanisms with Interdependent Valuations and Multidimensional Signals"; working paper (dated 07 Aug 2001).

McLean, R., and A. Postlewaite; 2002; "Informational Size and Incentive Compatibility"; *Econometrica* 70:2421-2453.

McLean, R., and A. Postlewaite; 2003; "Informational Size and Incentive Compatibility with Aggregate Uncertainty"; *Games and Economic Behavior* 45:410-433.

McLean, R., and A. Postlewaite; 2004; "Informational Size and Efficient Auctions"; *Review of Economic Studies* 71:809-827.

Plott, C.; 1994; “Market Architectures, Institutional Landscapes and Testbed Experiments”; *Economic Theory* 4:3-10.

Plott, C. 2000; “Markets as Information Gathering Tools”; *Southern Economic Journal* 67(1) (2000):1-15.

Plott, C., J. Wit, and W. Yang; 2003; "Parimutuel betting markets as information aggregation devices: experimental results"; *Economic Theory* 22:311-351.

Plott, C., and S. Sunder; 1988; “Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets"; *Econometrica* 56:1085-1118.

Ross, S.; 1987; Introduction to Probability and Statistics for Engineers and Scientists; John Wiley & Sons.

Roust, K.; 2004; "Informational Size and Behavior in an Information Aggregation Experiment"; working paper (dated 09 September 2004).

Würtz, A.; 1997; "A Universal Upper Bound on Power Functions"; UNSW Discussion Paper 97/17, University of New South Wales.

Appendix (Subject Instructions)

See the following pages for the five-page instruction packet received by subjects.

Name ______________________________________
Machine number ________________ Id ___________

INSTRUCTIONS

This is an experiment in the economics of market decision making. Various research foundations have provided funds for this research. If you follow the instructions carefully and make good decisions, you might earn a considerable amount of money, which will be paid to you in cash. The currency used is "francs". All francs will be converted to dollars at the end of the experiment, at a rate of ________ francs = 1 US cent.

You will be making decisions about participation in a type of "betting pool" process. The exercise will consist of several periods or "rounds". You will have the opportunity to participate or not according to your own decisions. Each round you will be given some cash, which you are free to keep or free to use to purchase tickets to the lottery. There will be six different types of ***event tickets*** for sale, labeled A-ticket, B-ticket, C-ticket, D-ticket, etc. The event tickets are associated with a lottery that pays according to the event that occurs at the end of the round. Each event ticket is associated with a particular event. The lottery payoff will be determined by the sale of all tickets plus a bonus payoff added to the payoff. The proportion of the winning event tickets that you purchased will determine your share of the lottery payoff.

Your payoff for a round will be determined by the following formula.

Part A Earnings = *your share of lottery (A) payoff*
Part B Earnings = *remaining cash + your share of lottery (B) payoff – period loan*
Round Earnings = Part A Earnings + Part B Earnings

Your share of lottery payoff =

$$\left[\frac{(\textit{Total francs from all ticket sales + bonus})}{(\textit{Total number of winning tickets sold})}\right] \textbf{[}\textit{number of winning tickets that you hold}\textbf{]}$$

Consider an example with three different events labeled A, B, and C and a bonus payoff of 27 francs added to the payoff. Suppose that the tickets are sold at a price of 1 franc each. Suppose further that you purchase no A-tickets, three B-tickets and one C-tickets. Suppose other people buy three A-tickets, six B-tickets (making a total of nine B-tickets sold, three to you and six to others), and five C-tickets. Thus ticket sales total to 18 considering all sales of all types. If tickets are priced at 1 franc each this gives a total of 18 francs contributed to the lottery payoff from ticket sales.

EVENT	1 tickets you purchased	2 total number of event tickets sold	3 lottery payoff sales + bonus	4 payoff per event ticket	5 Your lottery payoff if the event occurs
A-tickets	0	3	18 + 27 = 45	45/3 = 15	15 *x* 0 = **0**
B-tickets	3	9	18 + 27 = 45	45/9 = 5	5 *x* 3 = **15**
C-tickets	1	6	18 + 27 = 45	45/6 = 7.50	7.50 *x* 1 = **7.50**
total		**18**			

The table illustrates the payoff calculations. The sale of 18 tickets for all events produced revenue from ticket sales of 18 francs. All of the money from ticket sales is contributed to the payoff. In addition 27 francs is added as a bonus giving a total payoff of 45 francs. The total payoff is distributed to the holders of winning tickets. Since there were 3 A-tickets sold each of the three winning tickets would be paid 15 francs if A occurs. You hold none of them so your payoff would be zero. If event B occurs the 45 francs is distributed to the 9 winning tickets. Each winning ticket gets 5 francs and since you purchased 3 tickets your payoff from the lottery would be 15 francs if event B occurs. If the C event occurs you hold one of the six winning tickets so you would receive 7.50.

Each round you will be given a cash payment, which is yours to keep or yours to spend on event tickets. In addition to this cash payment you will be given a loan, which must be repaid at the end of the period. The loan is offered to give you the flexibility of investing in the lottery.

Hypothetical accounting example. If you are given a period payment of 50 francs plus a loan of 30 francs you will start with cash on hand of 80 francs. In the example you spend 4 francs for tickets leaving you with 76 francs on hand at the end of a period. The payoff you received from the lottery was 15 francs and you must repay the 30 franc loan, leaving you with a total income for the period of 61 francs.

period payment	50		
period loan	30		
initial cash on hand		50 + 30 = 80	
expenditures on tickets	4		
cash on hand before lottery payoff			80 – 4 = 76
payment from lottery (assuming event B occurs)			15
cash on hand plus payoff from winning tickets			76 + 15 = 91
loan repayment			- 30
Total period earnings			61

TICKET PRICES: The prices will be explained. Generally speaking the prices of tickets will increase with time during the period.

INFORMATION ABOUT EVENTS

SELECTION OF WINNING EVENTS. The winning event is chosen at random. Each of the six events is equally likely. Each period it is as if one of each letter is placed in an urn and one of them drawn at random. Which of the events wins will be announced after all ticket sales have closed.

illustration
It is as if one of each of the letters is placed in an urn and one letter is drawn as the state.

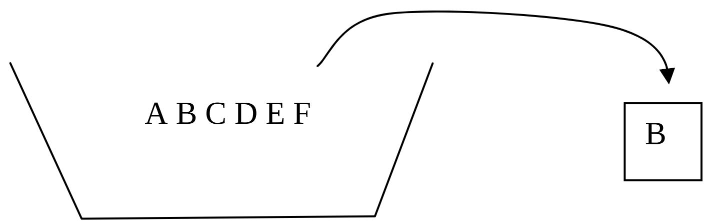

In this example the random draw is B. The chance that it is any particular letter is 1/6. It is important to note that the draw of a state is independent. The letter is returned to the urn after the draw, so the history of draws holds no implications of what future draws might be. It is just as likely that the draw every period is exactly the same letter as is any other sequence of draws.

PRIVATE CLUES. Each individual will be given clues about the winning event. The clues are determined independently for each individual by the following procedure. Once the winning event is determined a new urn is created with five letters of the winning event and two letters from each of the other events. Three random draws *WITH REPLACEMENT are* conducted and the results are given to a participant. Each participant is given the results of three independent draws.

The urn contains five of the winning state letters and ten letters of wrong states. The chance of the winning state drawn on any one draw is 1/3 for the winning state and 2/3 for a wrong state (2/15 for each of the five wrong states).

illustration when B is the winning state.

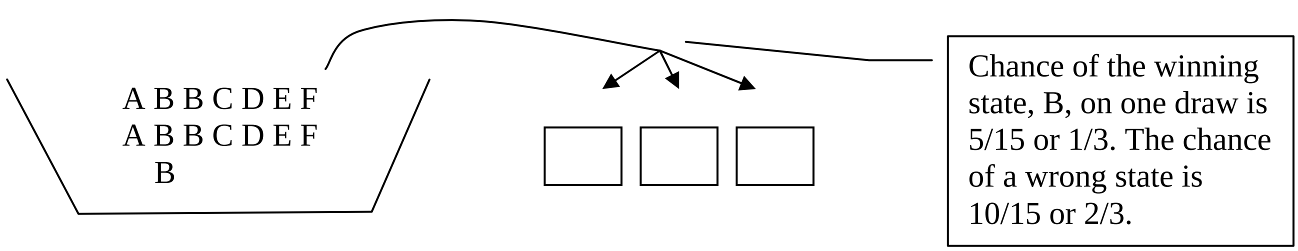

However, the three draws together contain more information. The probability calculations are below:

CHANCE THAT THE LETTER IS THE WINNING STATE BASED ON YOUR DRAWS				
	most frequent letter	second most frequent letter	third most frequent letter	each of the other letters
all three draws are the same letter	75.8%	4.8%	4.8%	4.8%
two of the three draws are the same letter	49.0%	19.6%	7.8%	7.8%
all three draws are different letters	23.8%	23.8%	23.8%	9.5%

The table contains the computations from classical models of draws, allowing for rounding errors. The chances that any particular letter is the winning state, given your clues, are listed in the table. If all three of your draws are the same letter then the chance that that letter is the winning state is 75.8%. Notice that you do not have perfect information since even if all of your draws are the same there is still a 24.2% chance that some other state is the winner.

If two of your draws are the same then the chance of that state being the winner is 49.0% (about 50:50) and the chance that the other letter drawn is the winner is 19.6% (about one chance in five). Any of the letters that is not one of these two has a chance of 7.8% (a little less than 1 in 10) but of course the chance that one of these four letters not represented in your clue is about 31.2% (7.8% times four).

If all of your draws are different letters then the chance that the winner is some particular letter among these three is 23.8% and each of the others not one of the three is 9.5%.

These numbers are the best representation of the information contained in your personal clues and in general these will be different from the clues and the information of other people. You should note, however, that the information distributed within the group as a whole is more than that of any one individual. Since each individual in the group has three draws, drawn independently, there are three times the number of individuals in the group draws in total. While this is a bigger sample and thus more information than any isolated individual has, the information distributed in the group is not sufficient to determine the winner with absolute certainty.

Structure of each round:

- Draw true state
- Draw private information for each individual

Part A (private purchases)
- Distribute private information
- Set clock and prices for part A
  1 minute
  Price = 100 francs (constant)
  Period Payment = 10,000 francs
  Lottery Bonus = 20,000 francs
    Remaining Balance is worthless.
    Only Lottery Winnings will be paid.
- Open part A buying
  Note that you are **un**able to see what other people are buying during part A
- Close part A buying
- Publish part A sales
- Calculate part A earnings (but do not distribute any information)

Part B (public purchases)
- Re-distribute same private information
- Set clock and prices for part B
  2 minutes
  Price = 100-220 francs (increasing)
  Period Payment = 25,000 francs
  Period Loan = 25,000 francs
  Lottery Bonus = 100,000 francs
    Remaining Balance will be paid
    Lottery Winnings will be paid
- Open part B buying
  Note that you **are** able to see what other people are buying during part B
- Close part B buying
- Calculate part B earnings

- Announce true state
- Distribute earnings information for both part A and part B
- Record earnings on Earnings Record Sheet

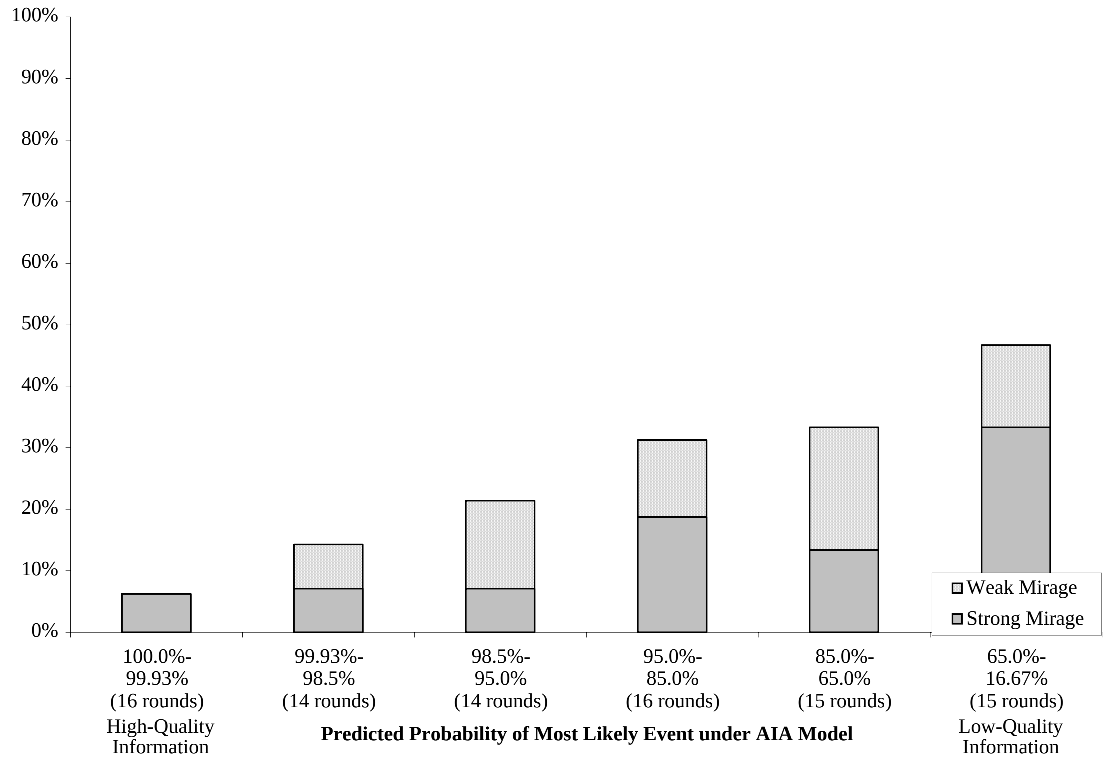

Figure 1. Mirages are more likely when lower quality information is available.

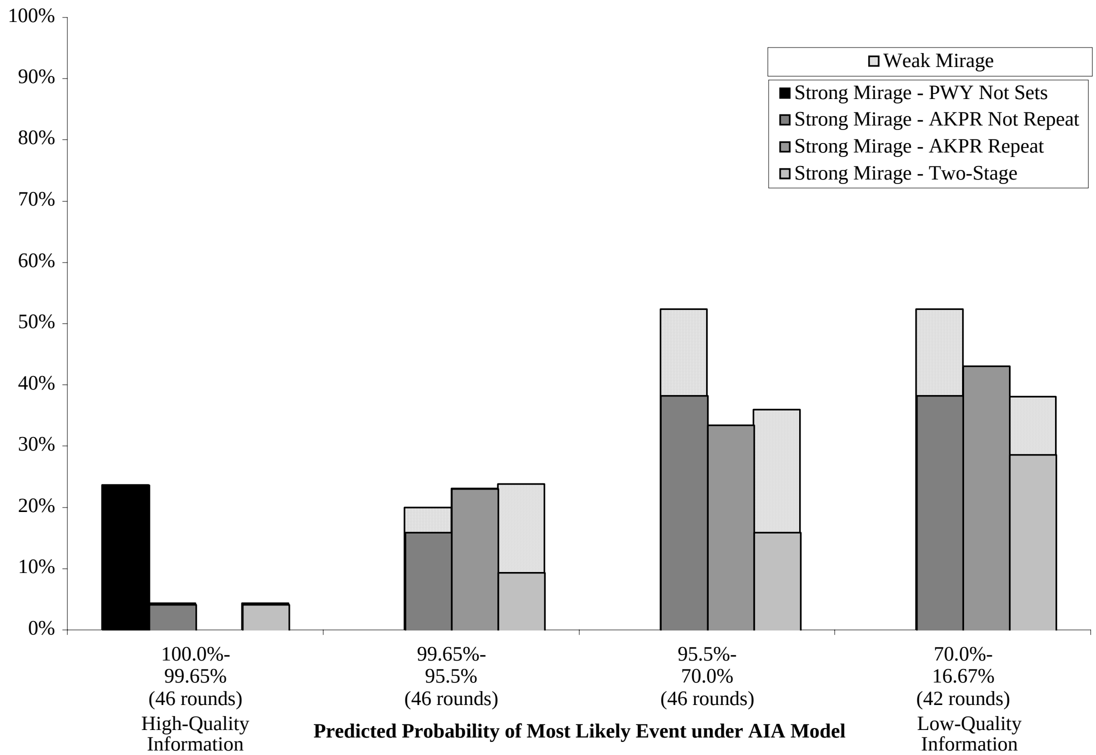

Figure 2. The Two-Stage Parimutuel IAM reduces the number of mirages, including a large decrease in strong mirages.

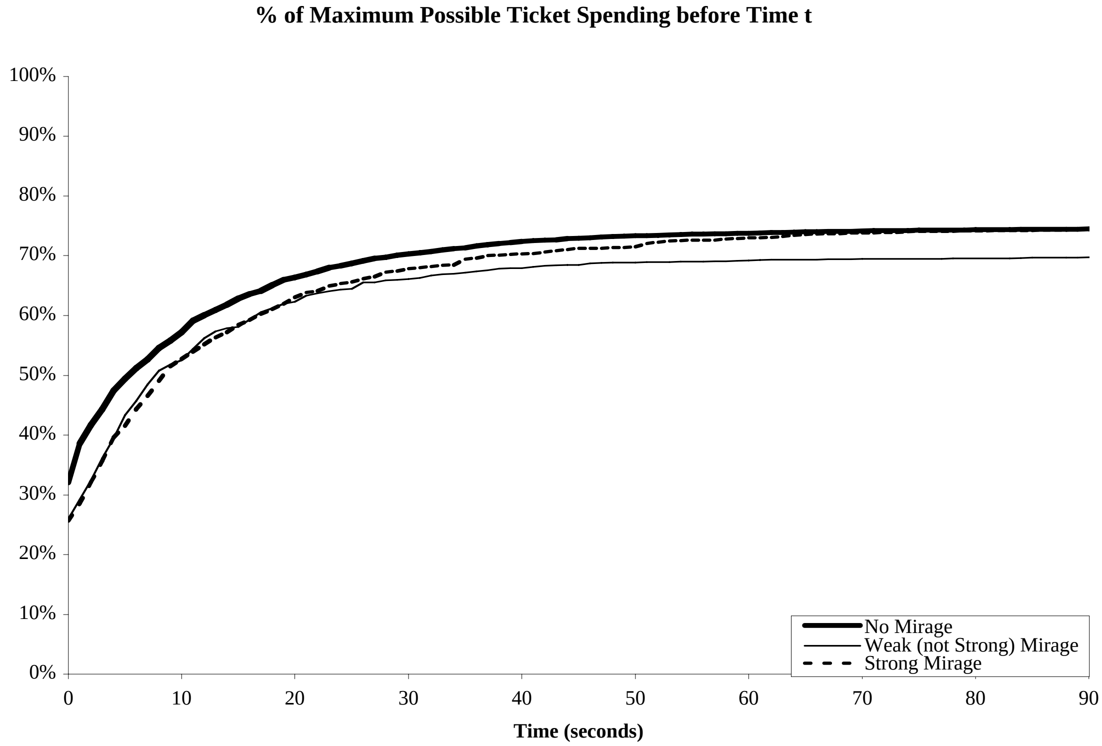

Figure 3. Purchases are made more slowly during mirages than during non-mirage periods. Weak mirages have less spending overall.

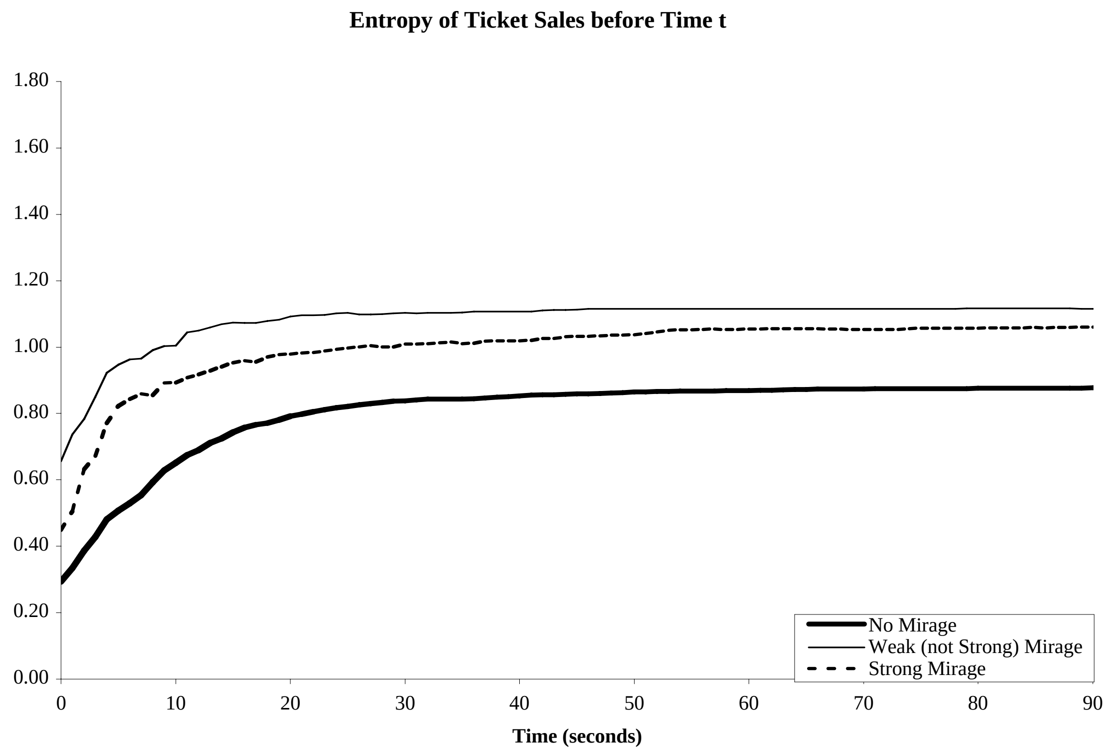

Figure 4. Purchases are more evenly spread across states during Mirage periods than during no-mirage periods.
