

The Double-Auction Gambling Market

An Experimental Examination

By KYLE W. HAMPTON*

ABSTRACT. I examine a new sports wagering mechanism that utilizes a double auction rather than the traditional wager matching bookmaker. Laboratory experiments were performed that modeled the two institutions in an attempt to compare the efficiency and volatility of the new mechanism versus the status quo. The new proposed double auction institution was less efficient and more volatile than the matching mechanism. However, the relative performance of the double auction mechanism was significantly improved by an increase in the size of the market.

I

Introduction

IN THIS PAPER, I COMPARE the information aggregating properties of new peer-to-peer gambling exchanges with more traditional bookmaking institutions. Economists have consistently found that sportsbooks predict the outcomes of sporting events with near perfect efficiency (Pankoff 1968; Woodland and Woodland 1994; Gray and Gray 1997). This finding is often marshaled to demonstrate the power of markets to utilize disparate and privately held information to make probabilistic predictions of uncertain outcomes.

However, this venerable institution is now facing competition from double-auction based “sports exchanges” that offer more flexibility in wagering and charge far less in commission than their traditional counterparts. As this institution grows in popularity, it offers a new

*The author is at the Department of Economics, University of Hawaii–Manoa. This research was conducted from spring 2002 to spring 2003 at the International Center for Economic Science at George Mason University. The author would like to thank Vernon Smith, Dave Porter, and Bart Wilson for their help and suggestions in producing this work.

opportunity to test the robustness of efficiency in gambling markets. Moreover, the double-auction structure of the new markets would seem to better parallel the financial instruments trading to which gambling markets are often compared.

It is difficult to measure the efficiency of market institutions with exclusive use of traditional empirical and theoretical economic analysis. However, by allowing us a greater degree of control of and knowledge about our market participants, economic experiments make it possible. Economic experiments have undergraduates earn real money by trading with one another in market simulations.

Smith (1962) and Plott and Sunder (1982) have utilized economic experiments to demonstrate that markets (specifically, double-auction markets) display a remarkable power to pool the myopic perspective of many participants into a singular measure. In most of these experiments, this measure is the price paid for some fictional asset. However, in the case of gambling markets, the aggregate measure takes the form of a prediction of an uncertain event. With recent findings suggesting information aggregating properties of markets are more sensitive to institutional and environmental specifications, I conducted an experiment to measure the predictive efficiency of both the traditional and new sports betting institutions.

With empirical data suggesting the predictive efficiency of sports betting markets and experimental data suggesting the efficiency of the double-auction institution, it seemed natural to proceed with the working hypothesis that the new institution, a double-auction sports book, would demonstrate the same powers of information aggregation. The results of my experiments have brought this assumption into question and raised new questions about the robustness of information aggregation in markets.

In the next section, I will describe the markets I am examining and the manner in which I model them in a controlled laboratory experiment. The following section will describe my findings and compare these data with predictions from the rational expectations model I hypothesize will describe the outcome. A fourth section will discuss the limitations of this particular set of experiments and consider changes for future research. The final section will summarize my conclusions.

II

Background

THIS RESEARCH IS THE FIRST to analyze sports wagering utilizing laboratory methodology. As a result, the literature offers little precedent in the modeling of these complicated decision-making environments. However, research on the predictive capabilities of markets includes both empirical analysis of traditional sportsbooks and experimental analysis of double-auction asset markets. From this, we should be able to extrapolate some expectation of efficiency in the markets we are studying.

The traditional sportsbook offers bettors the binary choice to wager on one team or another. Ideally, the money paid out to winners is less than or equal to the amount of money wagered. Any imbalance in the money wagered represents a risk being undertaken by the house. In order to balance the money wagered on each team, the house can offer different payouts for each or can specify a point-spread whereby some margin of victory or loss is considered in determining whether a wager wins or loses. For example, a contract whose associated team had a negative spread required that team to exceed the score of the second team by at least the value of the spread in order for the contract to win. If the difference simply equaled the spread, the contract tied. If it was less, the contract lost. Likewise, a contract whose associated team had a positive spread necessitated the team score at least the value of the spread less than the score of the other team. This method of balancing wagers allows a constant payoff of 2:1 for wagers won. This may account for these wagers being, by far, the most popular type. As a result, my research focuses exclusively on point-spread betting.

Setting a line does not guarantee that it will adequately balance the money being wagered. However, if an imbalance occurs, the line can be moved by some minimum increment. Typically, this increment is half a point. For example, if an imbalance has occurred where more money is being wagered on Team A than Team B, the line will move to where Team A's line is reduced by one-half of a point and Team B's line is increased by one-half of a point.

It is a distinctive characteristic of traditional bookmaking that wagers represent binding contracts. Wagers are accepted and paid

according to the line published when made. This is distinct from pari-mutuel gambling, in which the line being wagered is not determined until all wagers have been collected. Where pari-mutuel betting insulates the handicapper from exposure to losses, binding bets require that even if a line is moved, bets at the old line must still be paid. For their willingness to undertake this risk and provide the service of a market, the sportsbook charges a 10 percent commission to the bet winners. (Roxborough 1998)

Gray and Gray (1997) have pointed out that, unlike in traditional financial markets, the finite time horizon in sports wagers allows a clear point of comparison between the market-aggregated prediction and the actual outcomes of the sporting events. And as in the Iowa Electronic Markets,¹ the final line in the market for point-spread type wagers is an excellent predictor of the final score of games. This is found in markets for wagering on professional basketball (Brown and Sauer 1993) and football (Zuber et al. 1985; Pankoff 1968; Gray and Gray 1997; Gandar et al. 1988; Avery and Chevalier 1999; Dixon and Coles 1997). In each examination, market point-spreads at the end of the betting period were found to be unbiased estimators of the final point difference in the games.

New peer-to-peer sports exchanges would not, at first glance, give any indication that they deviate from traditional sportsbooks in their efficiency. The double-auction institution has long been utilized in the trading of securities and other homogenous commodities. Moreover, numerous experiments have demonstrated the efficiency of the institution in aggregating information over a wide range of environments (Smith 1962; Plott and Sunder 1988). Indeed, large-scale field experiments like the Iowa Stock Exchange (Forsythe et al. 1992) and the Euro 2000 (Schmidt and Werwatz 2001) futures markets have demonstrated the ability of double-auction institutions to efficiently predict the outcomes of exogenous uncertain events.

However, using a double-auction to broker point-spread wagers requires some modifications to the institution's structure. For each potential wager, the corresponding point-spread is the operative characteristic in the decision whether or not to "purchase" it. In this sense, we can think of dollars in the way we think of goods in a double-auction and the line as the price. A double-auction would need to

allow participants to bid and ask point-spreads for a particular number of dollars.

The auction proceeds with each subject being allowed to enter line bids on whatever stakes they choose. In the traditional case of buyers and sellers, we would use the terms *bid* and *ask* to denote the actions of a buyer or seller. However, in this case, there are not buyers or sellers in a traditional sense. Each individual is agreeing to a contract to pay or receive money contingent upon an independent outcome. A *bid* denotes this and specifies which side of the contract the wagerer is agreeing to.

Presumably, each subject has in mind an expected line. A deviation from this expectation in the wagers proposed by others would present a perceived profitable bet. If a participant, given his or her information, expects the result:

Team A: 75 points

Team B: 80 points

the participant would be willing to bet on Team A at line $>+5$ or Team B at line >-5 . Subjects would be allowed to enter bids on both sides of the market. In the example above, the subject could enter a bid to bet on Team A at line $+10$ and a bet on Team B at $+1$. These bids are placed in each teams' respective books. At this point, subjects are offered the ability to accept either side of each bid line wager. In our example above, a subject would take Team B at -10 or Team A at even odds. This represents a contractual wager between the two parties.

Subjects are also able to enter another line bid. In the example above, a subject could enter a line of Team A at line $+7$ or Team B at -2 . These would represent improvements and would represent the prevailing market lines. New bids that were not improvements would still be allowed, however. Participants would also be free to improve their bids at any time.

The double-auction sportsbook alters the role of the bookmaker. The bookmaker now functions as a mere broker. Rather than risk divining the demand for wagers on either side in setting the "price," the bookmaker merely coordinates contracts between bettors and functions as a third-party arbiter in determining a winner and assuring that the terms of the contract are met by each bettor. The brokerage

character of double-auction sportsbooks also increases the time in which wagers can be made. Traditionally, there is a limited betting period for each game in an effort by the house to minimize the imbalance in wagers created by an influx of new information. Given that the house is now insulated from the risks associated with changing information, wagers can be taken at any point prior and even during the contest. Moreover, with sports exchanges being freed from the costs of research necessary in setting a line, more events can be offered for wagering upon.²

This brokerage structure for wagers is currently being utilized by the burgeoning online gaming industry. In Britain, where online gambling is legal and licensed, Betfair brokers contracts between wagerers on a variety of sporting events. Included in their cadre of betting markets are double-auctions for team sports of the sort I have described here. The double-auction has also been made available at the offshore books WorldSportsExchange.com (2001) and Tradesports.com (2004).

Comparing the status quo institution with this new one in the laboratory will allow us to measure the efficiency of each with a degree of control not present in traditional analysis of sports betting markets. It also allows us to examine the information aggregating properties in the new high-information regimes made possible by the new institutional rules.

Defining efficiency in these markets is fairly simple. Sauer (1998) specifies the implication of efficiency in betting market as $PS = E(DP)$. PS is the point-spread determined by the market and DP is the actual difference in points that occurs in games. In other words, in an efficient market, the point-spread is an unbiased estimator of the scoring difference between the two teams. Another way to describe this efficiency is to demonstrate that no profit opportunity exists in the market. Traditional examinations of the efficiency of betting markets relied on post hoc measures of this efficiency wherein the closing lines of wagered games are compared to final scores to determine if there are any systematic biases. However, the laboratory allows a unique degree of control over the information provided to the bettors and thus the difficulty of the aggregation task the institution faces. In this way, these experiments differ from the election prediction markets

examined by Forsythe et al. (1992). We can examine not just the final lines but also the changes in the lines that take place over time as new information is provided to the participants. This should be enlightening given the ability of new sports exchanges to allow gambling in highly dynamic information regimes.³

III

Experimental Design and Procedures

AT THE CORE OF SPORTS BETTING is an event for which the bettors have differing probability beliefs about the outcome. Furthermore, the outcome of the game is typically exogenous and completely independent from the individual and aggregate actions of the bettors. In this experiment, subjects observe a game structured in this same way with two hypothetical teams each scoring points according to ambiguous probability distributions. Subjects place wagers according to their probabilistic beliefs as to the outcome of the game.

The game involves the periodic scoring of two hypothetical teams, Team A and Team B, each assigned a uniform distribution centered on a particular number. The size of each team's distribution will be the same for a single contest. The median of the distribution for each team may vary, however. Table 1 displays the teams' distribution characteristics for each of several parameter sets.

Scoring proceeds by drawing six times (one draw per period) from a team's assigned distribution and summing those draws over time to form a final score. This final score determines the payoff for contracts purchased by the subjects.

Splitting the process of determining the game outcome into several discrete parts allows us to create a game progression in which draws could be made sequentially to simulate scoring in a game. This allows prediction of the final score of the game with increasing accuracy as the game continues.

Each market consists of two stages, Stage 1 and Stage 2. An advantage of the double-auction-based sports exchange is the potential to allow betting during the game without additional cost. Modeling the game progression as a series of draws allowed us to present those draws sequentially to all the subjects, just as watching the game would

Table 1
Scoring Distributions

Parameter Sets	Market 1		Market 2		Market 3		Market 4		Market 5	
	A	B	A	B	A	B	A	B	A	B
Parameter Set 1										
<i>Distribution Median</i>	44	36	37	44	41	40	42	38	40	40
<i>Distribution Size</i>	7	7	11	11	11	11	11	11	11	11
<i>Expected Score</i>	264	216	222	264	246	240	252	228	240	240
Parameter Set 2										
<i>Distribution Median</i>	40	44	39	41	43	38	36	45	43	37
<i>Distribution Size</i>	7	7	11	11	11	11	11	11	11	11
<i>Expected Score</i>	240	264	234	246	258	228	216	270	258	222
Parameter Set 3										
<i>Distribution Median</i>	42	40	42	38	40	42	43	42	37	40
<i>Distribution Size</i>	7	7	11	11	11	11	11	11	11	11
<i>Expected Score</i>	252	240	252	228	240	252	258	252	222	260
Parameter Set 4										
<i>Distribution Median</i>	42	44	38	41	38	36	48	44	42	37
<i>Distribution Size</i>	7	7	11	11	11	11	11	11	11	11
<i>Expected Score</i>	252	264	228	246	228	216	288	264	252	222

present them a real-time accumulation of points. This allows us to structure an environment that models the traditional wagering decision (Stage 1) but also allows us to challenge the robustness of wagering institutions in alternate environments (Stage 2).

A. Stage 1

Stage 1 begins with a 30-second period during which each subject is presented several cards on which they can click to reveal a sample set of draws from the scoring distributions. These draws are presented as sample final scores determined in the same manner as the actual final score. Each draw is particular to the subject in order to simulate proprietary knowledge. Subjects are instructed as to the manner in which the final score is to be determined and the source of their draws. They are also instructed that their draws are unique to them and that other subjects may have different draws from the same distribution.

The remainder of Stage 1 consists of a three-minute betting period during which, on the basis of the information they were provided, subjects can purchase contracts. Each contract they purchase is displayed on the computer terminal as a record of their purchases.

B. Stage 2

The character of information presentation in this stage is that of a game in progress. Stage 2 immediately follows Stage 1 and begins with the first draw from the Team A and Team B distributions. This first draw is displayed to each subject in the form of a score. Five additional draws are made at 70-second intervals over the duration of Stage 2. These five timed intervals between each draw are referred to as *periods* in the instructions. The sixth and final draw occurs at the end of the Stage 2 after the time for the market had expired. Each new draw for the teams is added to the previous draws to form a cumulative score, as you would see in an actual sporting contest.

For each stage, subjects were allocated a budget of a fictional currency, francs, valued in real dollar terms at \$0.40 apiece. Subjects received four francs in Stage 1 and eight francs in Stage 2. A contract cost a single franc. Francs not spent during Stage 1 were not conveyed

to Stage 2. They were automatically banked. Likewise, the francs available in Stage 2 were not available for purchasing contracts until that stage had begun. Subjects were instructed that all unused francs would be automatically converted to earnings at the conclusion of Stage 2. Throughout this process of drawing, subjects were allowed to purchase contracts and see all contracts purchased in the market displayed on their screen. At the conclusion of Stage 2, each contract could win, lose, or tie. A contract that won would return double the purchase price of one franc, two francs.⁴ A contract that lost would return zero francs. A contract that tied would return the initial investment of a single franc.

Contracts are offered for each team, Team A or Team B, and an associated point-spread. For example, a contract may be described as Team B +8. This description details the terms by which the contract will pay off at the conclusion of Stage 2 when the final score is revealed. In this example, the holder of the contract would win if Team B's final score *plus eight points* totaled greater than Team A's final score.

At the conclusion of Stage 2, subjects are shown the final score and informed of the status of the contracts they purchased (i.e., win, lose, or tie). On the basis of these contracts, subjects' accounts are updated to reflect their provisional earnings in the experiment.

The institution for purchasing contracts appears in the computer interface during Stage 1 and Stage 2. The market is opened after the initial draws are revealed by the subject and closes after Stage 2 when the final score is revealed to the subjects. The two institutions we are interested in examining are the status quo matching mechanism and the double-auction alternative.

C. Matching

In the matching institution, subjects have the binary choice to purchase contracts for Team A or Team B with an associated point-spread. As the market progresses through Stage 1 and Stage 2, the line is updated in real time to resolve any imbalance in the contracts purchased. Subjects are not instructed as to how the line was determined.

The literature is unclear as to how much error we can expect from the house in setting the line (Gandar et al. 1988, 1998). As a convention for comparing the new mechanism with the status quo, I chose to give the status quo mechanism its best opportunity for success. The initial line represents the average of each draw given to the subjects. In other words, the opening line represents the aggregated information from each subject participating. The limited number of individuals in these experimental markets limits the volume of wagering we can expect. Also, the use of a discrete wagering currency, francs, limits our capability to set a threshold percentage difference at which we will alter the line.

In response to these constraints, I incorporated the following rule for the movement of the line: when the difference in the number of francs wagered on the two teams totals two, the line is adjusted by a single point. The two-franc difference is the smallest difference possible when using the discrete francs for purchasing contracts. In addition, the larger adjustment of an entire point rather than the customary half-point allows the line to adjust faster. The single-point adjustment is far easier to explain to subjects and is unlikely to harm accuracy, as the point totals in the game were larger than those typical in football or basketball games. The typical score for a single team was between 200 and 300 points.

D. Double-Auction

Like the matching institution, the double-auction interface appears on the screen while Stage 1 and Stage 2 are transpiring. The institution is split into two sections corresponding to Team A and Team B. Each section offers the subject the option to propose, or bid, a contract he or she would be willing to purchase. The section also allows the subject to accept a contract being bid by another subject in the market. Subjects are instructed that bidding a contract effectively offers a contract to other subjects. The contract being offered is simply the other side of the contract they have proposed. For example, bidding to purchase Team B at +7 is offering other subjects the opportunity to purchase a contract for Team A at -7. Because each contract is constrained to a single investment of one franc, subjects are agreeing

to a single wager contract between the two of them. In order to bid a contract, subjects first select the associated team's section on the interface⁵ and enter the spread they wish to bid. This bid is displayed for all subjects in the bid queue associated with the team. Additional bids on the same team contracts made by other subjects will also appear in the queue ordered according to terms. The bids with spreads most generous to those accepting the contract are placed at the front of the queue. For example, a bid for Team A +7 will be placed in front of a bid for Team A +9, as the latter contract has a greater chance of winning. Two bids with equal associated spreads are ordered according to the temporal order in which they were bid, with the earlier bid being placed above those after it.

At any time, a subject can choose to purchase a contract at the best current bid. In other words, if the Team A +7 bid is at the front of the Team A queue, the bid of Team B -7 will appear in Team B section with the option to purchase immediately. This same process is going on in the other team's queue, with the best bid being available with opposite terms on the other side. If a bid is accepted on either side, it represents a contract between the bidder and the purchaser with each having opposite terms. In effect, the contract is a wager between the two subjects. Upon the agreement to the contract, the associated bid is removed from the queue where it was displayed, and the next best bid in the queue is made available for immediate purchase by a subject.

Outside of actively choosing to purchase a currently bid contract, subjects would make an automatic contract purchase if they entered a bid whose terms were superior to the best bid currently being offered. For example, if the best Team A bid included the spread -5, and I bid to purchase Team B +3, I would also be willing to have purchased Team B +5. The contract will include the terms of the already bid contract. In the case of our example, the contract would be for Team B +5.

The subjects who have agreed to the contract each have a single franc deducted from their budget. Francs are only removed in the event of a contract. Subjects are free to make as many bids as they wish at any time on either team. However, when they have expended their last budgeted franc on a contract purchase, their remaining bids are removed from each team's queue.

The need to cancel bids due to changes in belief is important to the structure of the double-auction. Stage 2 involves constantly improving information that should alter the outcome predictions of market participants. One would expect that subjects would be very wary about submitting bids when new information could render those bids easily exploited late in the auction. In order to remedy this problem, the queues are cleared at the end of Stage 1 and at each period in Stage 2 when new information is displayed.

The experimental sessions were conducted at the laboratory of the International Center for Economic Science (ICES) at George Mason University. ICES is pioneering the use of experimental economics in the relatively new field of mechanism design to allow policymakers a controlled setting in which to analyze proposed changes in market institutions.

The experiment was administered using computer terminals connected via local area network. Subjects were recruited from undergraduate business and economics classes at George Mason University and were instructed they would be participating in an experiment in decision making. The experiment took place in a single laboratory with visually isolated terminals through which the subjects communicated exclusively.

Subjects were paid \$10 for arriving on time to both the practice session and the actual experiment. Though the subjects earned "francs," a fictional currency, during the course of the experiment, the students were advised that these francs would be converted to real dollars by a prespecified exchange rate. Based on this exchange rate, subjects could earn between zero and \$57.40 based on their decisions and the decisions of other experiment participants. The average actual subject earnings were \$21. All earnings and participation fees were paid in cash to the subjects immediately upon the conclusion of the experiment.

Experiments took place over four sessions. Each session consisted of 10 markets, five utilizing the matching institution and five utilizing the double-auction institution. In total, 20 markets utilized the double-auction institution for the trading of contracts; 17 used the matching institution.⁶ Subjects were trained in the mechanics of each institution through practice sessions and were selected for the data sessions on

the basis of their performance. In addition, quizzes were given to ensure subjects understood the instructions as given. Subject experience varies as each participated in several markets during a single session.

IV

Results

MY DESIGN ALLOWS US to analyze the efficiency of contracted spreads that occur in each market relative to subject expectations based on the information inputs we have provided. This allows a more accurate assessment of the market's performance throughout its operation, rather than simply at its close. Specifically, it allows us to analyze the ability of the market to adjust to new information and the consequent change in expectations. This ability to adjust quickly to changes in information is often discussed when describing the positive features of the double-auction institution.

Where equilibrium price is traditionally defined as the price at which the number of willing buyers equals the number of willing sellers, we can likewise define the equilibrium spread in these experimental markets as the spread where the number of subjects willing to wager on Team A equals the number of subjects willing to wager on Team B. Not surprisingly, this is exactly how traditional bookmakers describe their role in setting the line for a sporting event (Roxborough 1998).

In our experimental markets, the first step in determining an expected spread is to assume that each subject will average his or her draws to form a best unbiased estimate of the final score. The next step is to determine the equilibrium spread that would balance the wagers between either team, or the median of each subject's expected final score spread.

The in-game portion of the experiment involves five separate trading periods with new information presented to each subject at the outset. It is assumed that subject expectations of the final score outcome are adjusted as subjects receive more information on the final score through the sequential draws that take place. As a result, the expected spread should change. However, the draws should not

affect the expectations of future draws to take place. This assumes that the pregame wagering results in a common expectation as to the underlying distribution from which the scoring draws are being taken.

This allows the formation of a schedule of expected scores against which we can compare the contracted spreads in our experiment. This comparison will form the core of our analysis of the information-aggregating properties of our two institutions.

However, before we can make any conclusions based on this comparison, we should note an apparent serial correlation in contract pricing. In each market, there is a observed tendency for subjects to use the historical spreads in forming their expectations of the outcome. This results in contracted spreads that can be described as “sticky.” This appears to occur regardless of the institution.

This phenomenon is at its clearest when in the double-auction a subject purchases a contract at a spread highly divergent from that predicted. One of the supposed advantages of a double-auction institution should be its capacity to converge quickly toward the expected values. In many instances, this is exactly what takes place. For example, in the second double-auction market from the August data session, a contract is purchased for Team A -20 . Considering the expected spread value for this contract is Team $+44$, I make an assumption that this resulted from carelessness by the subject in reversing the signs when making a bid. However, the next contract spread reverts back to the general range of spreads.

Other times, however, it seems that subjects are using a divergent contract spread to reshape their expectations. There appears to be lingering effect when subjects contract for divergent spreads. This is manifested in a slow movement of contract spreads back toward the values expected.

This is a significant performance issue for the double-auction. The isolated mistakes of individuals participating in these markets seem to cascade by providing false information to other participants in the market. One of the touted benefits of the matching algorithm is that its very nature reduces the volatility of the market spreads. The spread is constrained to only move a single point at a time so it disciplines the

market and reduces the ability of a single contract (or a single subject) to seriously denigrate the efficiency of the market.

From the standpoint of examining the data, this serial correlation presents a more practical difficulty in measuring the information-aggregating performance of the two institutions. Parametric statistics such as summing the squared residuals depend upon an assumption of independence in contract spread we can no longer make.

Another nonparametric test of this divergence involves measuring the *mean absolute* deviation of the contract spreads. This measure was used by Plott and Sunder (1988) in their experiments measuring information aggregation in experimental securities markets. Their paper analyzes the actual prices and their deviation from the price predicted by various models of subject rationality. This is a very nice corollary for the comparison we want to make here.

The mean absolute deviation is defined as follows:

$$\frac{\sum_n |y_n^* - y_n|}{n},$$

where y^* is the expected spread for contract n and y is the actual contract spread.

Another measure of information aggregation in these markets is the percentage of price changes converging to the expected price, which allows us to get a better idea of the tendency of the market to encourage movement of the contracted spreads toward the values expected (Plott and Sunder 1988).

Table 2 displays the mean absolute deviation for each individual market and the percentage of convergent price changes. Again, following Plott and Sunder, I have chosen to analyze these data using the nonparametric Wilcoxon rank sum test.

For each parameter set, the matching institution does a significantly superior job of producing contract spreads that track the expected spread line.

One of the primary difficulties faced in measuring the performance of the double-auction is the highly divergent contracts that are purchased. These contracts can tend to inflate the mean absolute

Table 2
 Comparison of Contract Prices and Rational
 Expectations Prices

Experimental Markets			Percent Convergent Price Changes		Mean Absolute Deviation	
Parameter	Institution Order	Market	Match	Double	Match	Double
Set 1	Match Double	1	57.57	48.00	3.82	13.67
		2	47.62	48.00	1.89	7.83
		3	44.44	48.72	3.76	28.54
		4	47.62	53.33	2.87	35.59
		5		48.39		24.24
	Double Match	1	50.00	45.16	4.41	18.74
		2	48.65	51.85	1.68	7.94
		3	46.15	64.58	3.39	10.21
		4		50.23		17.42
		5		53.33		10.39
<i>Summary Statistics</i>			49		28	
Wilcox Rank Sum R_x			0.097 (Double)		0.000 (Match)	
Significance Level						
Set 2	Match Double	1	34.78	36.84	4.20	29.75
		2	53.85	54.90	8.47	16.68
		3	61.54	54.05	1.8	9.02
		4	57.69	37.93	3.10	6.32
		5	61.54	54.84	2.08	9.67
	Double Match	1	52.17	51.72	0.94	19.43
		2	81.25	57.89	7.75	14.10
		3	45.83	55.26	4.30	7.75
		4	34.78	38.46	4.68	7.17
		5	52.00	55.17	1.95	4.24
<i>Summary Statistics</i>			92		65	
Wilcox Rank Sum R_x			0.176 (Double)		0.001(Match)	
Significance Level						

deviation. Another measurement of the performance of the institution is the percentage of spread changes that are convergent on the expected spread.

In this respect, the double-auction performs much better in relation to the matching institution, with rank sum p -values of 0.097 and 0.176. This would indicate that though the double-auction does have the problem of producing divergent contract spreads, it is fairly adept at adjusting the contract spreads back toward the expected line.

The difficulty faced by the double-auction institution in consistently aggregating information runs counter to my stated hypotheses. It seems as though, in this environment, the ability of the double-auction to quickly adjust spreads when new information alters the expected outcome can also result in prices that diverge from the efficient prediction of that outcome. This raises many interesting questions about both my experimental design and the robustness of efficient aggregation in double-auction markets.

As a consequence of examining these results, I came to suspect that modeling this type of market using small numbers of people may be contributing to the difficulty faced by the double-auction in efficiently aggregating information. In order to determine the effect, if any, of the number of subjects on the performance of each mechanism, I ran a session with 16 participants, double the amount as before. Further, each of these subjects was given half the francs with which to purchase contracts (Stage 1: two francs; Stage 2: four francs) as they were given in the original experiment with the value of each franc being doubled. No other changes are proposed for the experiment.

Table 3 shows the results of these experiments. You can see that, compared to the original double-auction markets, the larger markets track the expected spread much more closely.

Once again, the matching institution significantly outperforms the double-auction institution. However, the absolute size of the mean absolute deviations in the double-auction institution has been reduced from an average of 17.46 for Parameter Set 1 and 12.41 for Parameter Set 2 down to 4.66 for Parameter Set 3 and the doubled subject pool.

Table 3
 Comparison of Contract Prices and Rational Expectations
 Prices—Large Group

Experimental Markets		Percent Convergent Price Changes		Mean Absolute Deviation		
		Match	Double	Match	Double	
Parameter	Market					
		1	42.86	44.44	3.65	4.97
	Set 3, Double	2	55.17	50.00	5.62	3.72
	Set 4, Match	3	52.17	38.64	3.01	5.45 ^a
		4	52.38	55.56	1.68	5.44
	5	45.33	50.00	2.45	3.70	
<i>Summary Statistics</i>			25		20	
Wilcoxon Rank Sum R_x			0.345 (Match)		0.075 (Match)	
Significance Level						

The information-aggregating properties of the double-auction are definitely improved by the addition of more people, more time, and fewer francs. The institution still performs worse relative to the status quo, but the gap shrinks. This indicates that future experiments with larger numbers of subjects may be successful in demonstrating, in the lab, the power of the double-auction to aggregate information in wagering markets.

V

Summary and Conclusions

THE IMPROVEMENT IN EFFICIENCY that occurs when the number of participants is increased is encouraging, especially in that experiments modeling these types of markets remain, for the most part, uncharted waters for economics. However, it is probably helpful to consider the difficulties that this increased number of subjects ameliorates.

Unlike traditional double-auction experiments, where the induced valuations reveal a very clear relationship between the price paid for a

unit and the resulting profit or loss, the information (the effective valuation for a good of ambiguous value) provided to the subject is only a clue as to the uncertain profit a particular contract will bring. Unlike the case of traditional double-auctions, in which the lessons of a profit or loss is clear, a loss or win in the betting markets is not as clear a signal that the decision one made to purchase a contract was the right one.

As a result, subjects who utilize faulty heuristics to deduce the probability of a final outcome may get positive feedback from their actions and continue to use that faulty heuristic. Likewise, subjects who utilize rational expectations and deduce unbiased estimates of the equilibrium spread may still lose due to a final score draw that diverges from their expectations. As a result, the subjects may give up on their means of deducing outcomes.


This conclusion is supported, at least anecdotally, by observing the actions of subjects in the experiments. Often, subjects would raise their hand in frustration, having negotiated very advantageous contracts only to see their contracts lose due to a peculiar draw for the final scores. Other times, subjects refused to even look at their information prior to the experiment, depending instead on luck to win contracts.

This phenomenon cast doubt on the robustness of the double-auction institution when faced with assets of uncertain and probabilistic value. Experiments involving securities with ambiguous returns have already observed the tendency to “bubble” as subjects speculate upon the returns from the asset (Smith et al. 1988). This divergence from the intertemporal equilibrium occurs in a very clear pattern of overvaluation for an asset. In the case of our experiments, there is no single asset being traded, as each contract represents the assumption of two different outcomes on the part of the subjects. It seems possible that the same type of expectation formation that is occurring in asset markets is taking place in our experimental wagering markets but is complicated by the dual ownership of a contract.

One important feature of the double-auction mechanism is the ability to provide efficient outcomes despite the diverse expectations of the participants. Whatever mistakes in expectation are made by the subjects, one would not expect these mistakes to be systematic enough to seriously undermine efficiency. However, in the case of a smaller market, particularly a market like mine in which subjects are

Figure 1

Relationship of Average Contracted Price to Expected Price (Bid Contracts)



prone to forming expectations based on the actions of others in market, the low volume of trades can result in some mistaken expectations finding their way into the spreads being contracted.

The gaps between spread bids for each team never converged tightly, allowing the diverse expectations in the market to manifest themselves in the form of divergent prices. One critical finding from examining the decisions of individual subjects is the relative unwillingness of some subjects to submit bids into the queues, as opposed to simply accepting those being offered. This would not be interesting except that this behavior turns out to be correlated with divergent contract spreads in the experiment. Figure 1 shows the advantageous terms of contract spreads garnered by those who had their bids accepted (as opposed to those who simply took offered bids).

Besides increasing the number of participants, the later sessions increased the value of each franc to the subjects and the time they had to spend them in the market. This seemed to be an effective means of reducing the cognitive costs for the subjects and induces them to wait for the queues to develop and more advantageous spreads to be bid before accepting the current offered spreads.

The reduced number of francs also reduces the path-dependent information cascades that successive contract purchases from a single subject might cause. Each of these factors would explain the substantive improvement of the double-auction performance in the larger experiments I ran and, by extension, larger volume markets outside the lab.

One final difficulty faced in these smaller markets is the seeming heterogeneity in subjects' understanding of the environment. This is most apparent in the bimodality of subject earnings. Subjects who do well, do well consistently, and vice versa. The disparity tends to shrink in the later markets, suggesting a marked learning effect despite an entire day of training and selective admission into the experiment.

One explanation for this disparity in performance may be the difficulty of internalizing the manner in which contracts are defined. The convention of defining an underdog spread with a "+" seems, in many ways, counterintuitive. Several subjects had difficulty understanding it. Though the instructions and interface do not use the language of wagering at all, several subjects experienced in sports gambling picked up on the connection right away.

This confusion seems to be accountable for many of the especially divergent contract spreads that occur in the double-auction. Often, subjects are trading at the exact opposite contract spread that is predicted by the rational expectations model. The unfortunate consequence of this action is that the false spread is published and forms a focal point for additional contract spreads to be bid.

Despite these difficulties, the results of this experiment demonstrate the ability of double-auction gambling markets to efficiently aggregate dispersed information and provide unbiased predictions of uncertain outcomes. These results hint at the potential benefits of using wagering markets in both industry and government to capture the wisdom of interested and dispersed parties in order to better anticipate the future.

Notes

1. The market operates by allowing participants to purchase portfolios of candidates for \$1. In the case of the 1992 election, each portfolio included four shares: George H. W. Bush, Bill Clinton, H. Ross Perot, and Other. After the election was complete, each share would pay, as a percentage of a dollar, the

percentage of the popular vote each candidate received. In 1992, the market closed on the eve of election day and yielded a prediction of the outcome better than any poll conducted at the same time (Forsythe et al. 1992).

2. Tradesports.com has offered wagering on political events around the world, including the date of capture for Saddam Hussein and Osama bin Laden.

3. Tradesports is configured to allow bettors to place wagers while the contest is ongoing. Traditional markets have limited in-game betting to half-time lines available only until the second half gets under way.

4. In order to induce activity in the markets, a 20 percent bonus was added to the earnings for winning contracts. This was meant to emulate Sauer's (1998) utility of winning above and beyond the consequent earnings.

5. The double-auction institution displays two different queues, each housing the bids for the two different teams. Next to these queues are the boxes for bidding accepting contract spreads for that team.

6. Software difficulties resulted in lost or unreliable data for three of the markets.

7. During training, the subjects went through an experiment iteration and were provided the instructions prior to the corresponding portion of the experiment.

References

- Avery, Christopher, and Chevalier Judy. (1999). "Identifying Investor Sentiment from Price Paths: The Case of Football Betting." *Journal of Business* 72(4): 493–521.
- Brown, William, and Sauer Raymond. (1993). "Fundamentals of Noise: Evidence from the Professional Basketball Betting Market." *Journal of Finance* 48(4): 1193–1209.
- Dixon, Mark J., and Stuart G. Coles. (1997). "Modelling Association Football Scores and Inefficiencies in the Football Betting Market." *Applied Statistics* 46(2): 265–280.
- Forsythe, R., F. Nelson, G. Neumann, and J. Wright. (1992). "Anatomy of an Experimental Political Stock Market." *American Economic Review* 82(5): 1142–1161.
- Gandar, John, Richard Zuber, Thomas O'Brien, and Ben Russo. (1988). "Testing Rationality in the Point Spread Betting Market." *Journal of Finance* 43(4): 995–1008.
- . (1998). "Informed Traders in the Betting Market for Professional Basketball Games." *Journal of Finance* 53(4): 995–1007.
- Gray, Phillip K., and Stephen F. Gray. (1997). "Testing Market Efficiency: Evidence from the NFL Sports Betting Market." *Journal of Finance* 52(4): 1725–1737.

- Pankoff, Lyn D. (1968). "Market Efficiency and Football Betting." *Journal of Business* 41(2): 203–214.
- Plott, Charles P., and Sunder Shyam. (1982). "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational Expectations Models." *Journal of Political Economy* 90(4): 663–698.
- . (1988). "Rational Expectations an the Aggregation of Diverse Information in Laboratory Security Markets." *Econometrica* 56(5): 1085–1118.
- Roxborough, Michael, and Mike Rhoden. (1998). *Sports Book Management*. Las Vegas, NV: Las Vegas Sports Consultants.
- Sauer, Raymond. (1998). "Economics of Wagering Markets." *Journal of Economic Literature* 36(4): 2021–2064.
- Schmidt, Carsten, and Werwatz Axel. (2001). "How Accurately Do Markets Predict the Outcome of an Event? The Euro 2000 Soccer Championships Experiment." Discussion Papers on Strategic Interaction 2002-9, Max Planck Institut for Research into Economic Systems, Jena, Germany.
- Smith, Vernon L. (1962). "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy* 70(2): 111–137.
- Smith, Vernon L., Gerry L. Suchanek, and Arlington Williams. (1988). "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica* 56(5): 1119–1151.
- Tradesports.com. (2004). <http://www.tradesports.com>
- Woodland, Linda, and Bill M. Woodland. (1994). "Market Efficiency and the Favorite Longshot Bias: The Baseball Betting Market." *Journal of Finance* 49(1): 269–279.
- World Sports Exchange. (2001). <http://www.wsex.com>
- Zuber, Richard, John M. Gandar, and Benny D. Bowers. (1985). "Beating the Spread: Testing the Efficiency of the Gambling Market for National Football League Games." *Journal of Finance* 93(4): 800–806.

Appendix: Experiment Instructions⁷

Instructions—Stage 1

This is an experiment in the economics of decision making. Various research foundations have provided funds for this research. The instructions are simple, and if you follow them carefully and make good decisions, you may earn a considerable amount of money that will be paid to you in CASH at the end of the experiment. Your earnings will be determined partly by your decisions and partly by the

decisions of others. If you have questions at any time while reading the instructions, please raise your hand and a lab monitor will assist you.

Please Note: At the end of the instructions, you will be asked to answer some questions to verify that you understand the experiment.

In this experiment, you will be able to purchase **contracts** that can earn you profit. This profit will depend on your ability to predict the difference of two scores, **Score A** and **Score B**. For example, Score A may be **199** and Score B may be **238**. These scores will be determined randomly and independently of your actions in the experiment.

Furthermore, Score A and Score B are determined independently from one another by drawing six times from each of two different number lines. The number lines will be the same length but may be centered differently. For example, Score A may be drawn from this number line:




while Score B is drawn from this one:




Each spans 10 possible numbers. However the second number line is centered 7 points higher than the first. **Please Note:** The number lines given above are merely examples.

The **Final Score** will be determined by drawing six times from each number line and adding up the draws. Continuing with our example, the following draws may be made for Final Score A: 35, 33, 31, 37, 33, 38.



This results in a Final Score A of **208** [35 + 33 + 31 + 37 + 33 + 38 = 208].

Final Score B may be determined by the sum of these draws: 37, 46, 42, 44, 40, 39.



This results in a Final Score B of **248** [$37 + 46 + 42 + 44 + 40 + 39 = 248$].

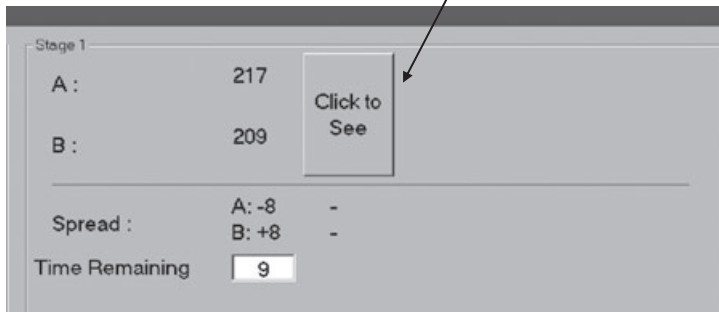
A: 208 | B: 248 represents the **final score**.

This experiment will consist of 5 **rounds**. Each round will involve a final score like the one described above. The final score *will not be revealed until the round is over*. Each round consists of Stage 1 and Stage 2.

- In **Stage 1**, you will be given information to help you predict the Final Score.
- In **Stage 2**, you will actually see the draws take place that will be summed to form the final score.

Stage 1

The experiment begins in **Stage 1** with the appearance of the box shown below. You will be shown several cards and be given 30 seconds to click each one, revealing a score underneath. These scores are determined in the exact same way as the final score will be determined later in the experiment and use the exact same number lines as will be used for the final score. You can use this information to predict the final score.



- Each score shown represents a completely independent set of draws from the same number lines that will be used to determine the final score.
- The scores shown are not a final score. They are simply determined in the same way as the final score and can be used as information to predict the final score.

Earning Money

At the end of the draw period, you will be able to purchase contracts.

You will be given a budget of **francs** to spend purchasing contracts. Each contract costs one franc. A contract can pay you a profit. The payoff of the contract depends on the final score reported at the end of the round. Each contract loses, wins, or ties according to rules we will discuss later.

- If your contract **loses**, you simply lose the franc you spent to purchase the contract.
- If your contract **ties**, you will be refunded the franc you spent purchase to purchase it.
- If your contract **wins**, you get your franc refunded back with a bonus. This bonus is 1.2 francs.

For example, if you won a contract, you would get 1 franc (refund) + 1.2 francs (bonus) = **2.2 francs**.

At the end of each round, after the final score is reported, your total earnings will include *francs won on contracts* and any *francs that you did not spend*. These francs will be reported in an earnings box at the bottom of your screen. You will be told the exchange for francs to U.S. dollars before the experiment begins.

Winning and Losing Contracts

Based on the information provided under the cards, you can purchase either an **A** contract or a **B** contract. Each of the two contracts will include a **spread**. A spread describes a difference in score A and score B. For example, the contracts may be

Contract A: +24

Contract B: -24

Each contract will win or lose depending on its spread compared with the Final Score determined at the end of the round. The following describes the conditions where each contract would win. Please pay close attention to the conditions for winning and losing given a particular spread.

Suppose you purchase contract B: -24. That contract:

- **Wins** if the final score B is more than 24 points higher than Final Score A.
- **Loses** if B is less than 24 points higher than A.
- **Loses** if A is higher than B.
- **Ties** if B is exactly 24 points higher than A.

For example, if you have purchased **B: -24** and the final score is:

- **A: 190 | B: 228**, the contract **wins** because final score B is **38** points higher than final score A.
- **A: 190 | B: 205**, the contract **loses** because final score B is only **15** points higher than final score A.
- **A: 190 | B: 181**, the contract **loses** because final score B is **less** than final score A.
- **A: 190 | B: 214**, the contract **ties** because final score B is *exactly* **24** points higher than final score A.

Now suppose you purchase contract A: +24. Your contract:

- **Wins** if the Final Score A is less than 24 points lower than Final Score B.
- **Wins** if Final Score A is higher than Final Score B.
- **Loses** if B is more than 24 points higher than A.
- **Ties** if B is exactly 24 points higher than A.


For example, if you have purchased **A: +24** and the final score is:

- **A: 190 | B: 205**, the contract **wins** because final score A is only **15** points less than final score B.
- **A: 190 | B: 181**, the contract **wins** because final score A is **greater** than final score B.
- **A: 190 | B: 228**, the contract **loses** because final score A is **38** points less than final score A.
- **A: 190 | B: 214**, the contract **ties** because final score A is *exactly* **24** points lower than final score A.

Please Note: The spread for one of the offered contracts is the simply the opposite of the spread offered for the other. In other words, of the contracts offered, one will win and one will lose or both will tie.


Purchasing Contracts

In this institution, each participant will be allowed to purchase multiple contracts. On the left side of your screen, you will see the market where you purchase contracts. You will be able to choose the A contract or the B contract at the spreads shown by clicking the **Take** button beside them. For example, in the market below you could choose to purchase contract A: +19 or B: -19.



Note: As contract purchases take place, the contract spreads may change to balance the number of A and B contracts sold.

After purchasing a contract, you will see it appear in the box on the right as a purchased contract.



You are allowed to purchase as many contracts as you wish before the timer runs out. Again, if you have uninvested francs at the end of the purchasing period, you are allowed to keep them.

This will conclude Stage 1 of the round. Before we continue to Stage 2, we will practice what we have learned up to this point. If you have any questions, please raise your hand. Otherwise, please **complete the quiz on your computer screen**.

Instructions—Stage 2

After Stage 1 ends, we will move immediately into Stage 2.

The screenshot shows a software interface for Stage 2. On the left, there are several input fields: 'Period' set to 3 of 5, 'Event A' set to 131, 'Event B' set to 134, 'Francs Remaining' set to 27, 'Contract Earnings' set to -, and 'Total' set to -. On the right, there is a table with the following data:

Buy #	Event	Spread	Earnings	Result	Take/Bid
1	B	+11			Bid
2	A	-12			Take
3	B	+4			Take

At the top right of the interface, 'Period Time Remaining' is set to 2.

In Stage 2, this box will appear in the lower right-hand corner of your screen. You will be given a new budget of francs with which to purchase contracts based on your prediction of the final scores for A and B. The rules for purchasing and winning contracts are exactly the same. However, you will now be provided additional information about the final score as Stage 2 progresses.

The final score is calculated, as described before, by drawing six times each from the A and B number lines. During Stage 2, *these draws will take place slowly while you trade*.

Stage 2 consists of 5 **periods**. In the Event A and Event B boxes shown above, Score A and Score B will be increased incrementally, or **updated**, at the start of each period until the end of Stage 2 (and the

end of the round). These updates represent each new draw from each number line being added to the total score currently displayed. The sixth and final draws are added at the end of the last period to form the Final Scores.

Imagine the Final Score in the example above is **A:269 | B: 264**. This final score could be revealed during Stage 2 as follows:

Period	Event A		Event B	
	Draw	Total Score	Draw	Total Score
1	43	43	47	47
2	41	84	44	91
3	47	131	43	134
4	42	173	41	175
5	49	222	42	217
Final	47	269	47	264

The shaded row represents the period shown in the example above.

You will be able to purchase contracts during Stage 2 while the scores are updating. You may purchase contracts as early or as late in the Stage as you wish. You purchase contracts in the same manner as you purchase them in Stage 1.

Please note: The queues will be cleared at the end of each period. You will need to enter new bids at that time.

When the final scores are revealed, you are paid your earnings based on this final score and the contracts you purchased in both Stage 1 and Stage 2. Your record of contracts purchased will show if each contract won, tied, or lost. In order to help you understand

why you won or lost on a contract, you will be told how much you won or lost by.

Here is an example of how your recorded contracts may appear after the round has ended. Please look at the **Result** column and confirm you understand how contracts are won and lost.

Stage 1

A : Click to See Click to See Click to See

B :

Spread : - - -

Time Remaining	0	Buy #	Event	Spread	Earnings	Result	Take/Bid
Francs Remaining	16	1	A	-4		2 W by 1	Bid
Contract Earnings	5	2	B	+5		1 TIE	Bid
Total	21	3	B	+8		2 W by 3	Take
		4	B	+3		0 L by 2	Take

Stage 2

Period of 5 **Period Time Remaining**

Event A

Event B

Francs Remaining	27	Buy #	Event	Spread	Earnings	Result	Take/Bid
Contract Earnings	2	1	B	+11		2 W by 6	Bid
Total	29	2	B	+2		0 L by 3	Take
		3	A	-12		0 L by 7	Take

Continue
Round
Total Earnings

The cumulative **Earnings** from each contract are added to your earnings total shown at the bottom right on your screen. In the example shown, the earnings for the round were 50 francs. The Total Earnings box represents these earnings and earnings from previous rounds. Please carefully examine the round's results. If you have any question, please raise your hand and someone will gladly help you. When you understand why you won and/or lost contracts, and are ready to continue to a new round, you click the **Continue** button at the bottom of the screen.

Each round is an entirely different game. Final scores are determined from new number lines of different size and location.

We will now finish our practice round by proceeding to Stage 2. If you have any questions, please raise your hand. Otherwise, please sit quietly until everyone has finished the instructions.

Purchasing Contracts—Another Institution


In this institution, each contract you purchase requires another participant to purchase the opposite contract. In other words, if you wish to purchase the contract A: +24, another participant would have to enter an **agreement** to purchase the B: -24 contract. Their contract is simply the opposite of your own. If your contract wins, their contract loses. If your contract loses, their contract wins. If you tie, they also tie.

While your purchase price per contract stays constant at one franc apiece, the spread is variable. You seek agreements to contract in two ways:

- You can **bid** to take one side at a particular spread and have someone accept your offer to purchase the opposite contract.
- You can **take** a contract that has the opposite spread of another contract being offered by someone else in the market.

Bidding

The market you see here will appear on the left side of your screen during the experiment. In order to bid, you will choose which type contract you want, A or B.




Imagine that you are willing to purchase an A contract at the spread +24. In other words, you would win as long as Final Score B is not 24 points more than Final Score A. In order for you to purchase A: +24, someone else would have to agree to take B: -24. In other words, they would have to be willing to take a contract that would only win if Final Score B was over 24 points more than A.

In order to match up with others, you could propose a trade to others. This is called **bidding**.

If you want to purchase B, you would bid at the bottom box where it says B. If you want to purchase A, you would use the top box where it says A. The box allows you to enter the spread for that contract you are bidding.

For example, if you wish to purchase A: +24, you would enter "+24" in the box next to A. When you click the **Bid** button, your bid will appear in the A **queue** to the left of the box. This **queue** appears to everyone else in the market and shows all of the bids made for the A contract.



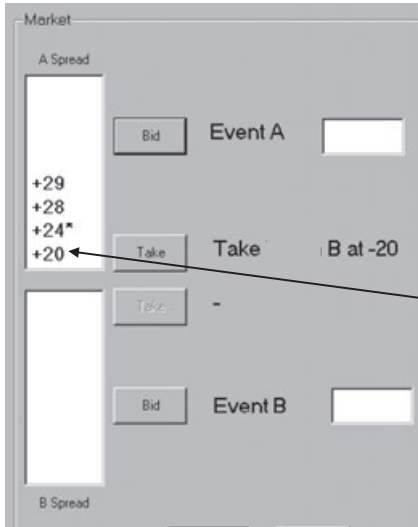
Other people in the experiment may also want to purchase an A contract and can bid as well. Continuing with our example, you can see bids that have been made by other people along with your A: +24 bid. The bids are ordered according to their spread. The lowest spread numbers are placed closest to the middle of the market. For the A queue, this means the bottom.

At any time during Stage 1, any person can choose to purchase the other side of the contracts that have been bid. This is called a **Take**. When you take, you purchase the best available bid in the queue. In other words, the contract spread with the best chance of winning.

Here, someone is offering A: +15. This means that anyone can **Take** B: -15. This is better than taking the B: -24 that your bid of A: +24 requires.

- For B: -24 to win, Final Score B has to be over 24 points greater than Final Score A.
- For B: -15 to win, Final Score B only has to 15 points greater than Final Score A.

If you wanted to Take B: -15, you would click the **Take** button next to it.




Once someone Takes B: -15, that person and the person who bid A: +15 have each purchased a contract. The A: +15 bid is removed from the queue and the next best bid, A: +20 moves to the front.

Buy Record

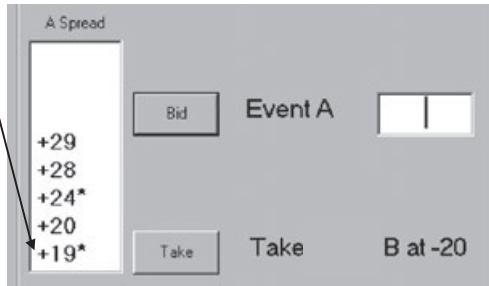
Period		
A	-	
B	-15	

At the bottom of the market, you will see a record of each contract purchased and the spread at which they were purchased.

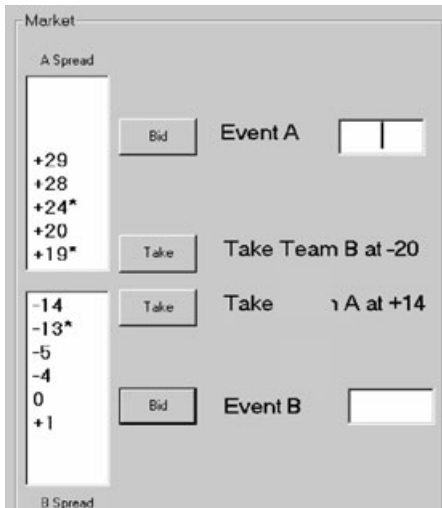

If no one is willing to **Take** your bid of A: +24, you may need to enter an improved bid. For example, you may enter A: +19.



This contract would have reduced chance of winning but an improved chance of being taken.



Of course, you can also Bid to purchase B contracts in the box at the bottom half of the screen. For example, you may enter a Bid to purchase B: -13.



This contract will win only if B is over 13 points greater than A. You are looking for someone to **Take** A: +13. In the example shown, your bid would need to wait until someone took A: +14 to remove B: -14 from the front of the queue.

These are the basics of bidding for and taking contracts in this experiment. If you have any questions, please raise your hand. Your understanding of these markets will determine how much you money you make.

Here are some additional rules and tips for this market:

- You can enter bids or take bids at any time during Stage 1 until time runs out or you run out of francs. You are not charged a franc to make a bid.
- You only pay a franc when you have purchased a contract. When you run out of francs, any bids you have in the queues will be removed.
- You are free to enter the same spread in your bid as another bid in the queue. When two bids are the same spread, the bid entered first will be taken first.
- Your own bids in the queue will be marked with an asterisk. You cannot take your own bids.
- If you enter a bid where you could have simply taken a bid and received a better contract, you will automatically take the bid. For example, if B: -15 is available to be taken, and you enter a bid, B: -19, you will take B: -15 instead.
- You can bid for or take contracts for both A and B. For example, if you believe Final Score B will be 10 points more than Final Score A, you may want contract B: -5 or contract A: +15.
- Unspent francs are added to your earnings.

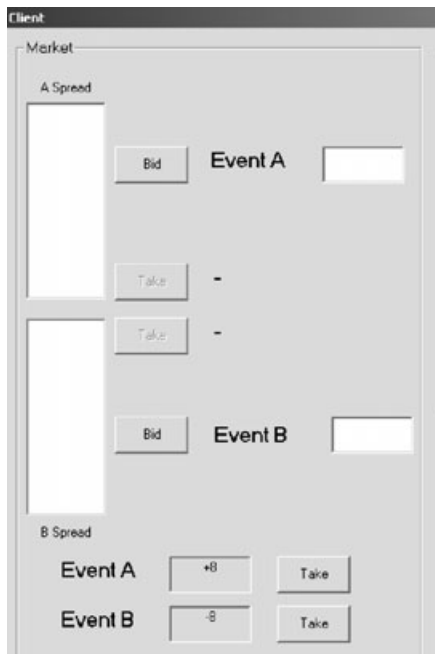
The contracts you purchased will be listed on the right side of your screen. For example, if your bid of A: +19 is accepted, you will see the contract listed as you see below.

Time Remaining	343	Buy#	Event	Spread	Earnings	Result	Take/Bid
Francs Remaining	19	1	A	+19			Bid
Contract Earnings	-						
Total	-						

This institution will be used in both stages. We will go through Stage 1 now to give you some practice. If you have any questions, please raise your hand. Otherwise, please sit quietly and wait for others to finish.

Dual Institutions

In this session, you will have the option to purchase contracts in either of the two institutions described above. Each institution will appear in the box on the left side of your screen.



You can use either institution at any time during the experiment to purchase contracts. For example, you may choose to spend your first franc purchasing a contract in the top institution where bids are submitted and accepted and then spend your second franc purchasing a contract in the second institution. You can use either institution as much or as little as you like during the round.

You still have a single budget of francs. A purchase of a contract using either institution will decrease your budget by one franc.

If your budget is depleted to zero francs by a contract purchase in either institution, any outstanding bids you have in the first institution will be removed automatically.

As described above, winning contracts are paid a bonus. This bonus will be changed during the experiment. The experimenter will notify you of the bonus before you begin.