

## TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

*Ricard Gil†\* and Steven D. Levitt‡*

†*University of California-Santa Cruz*

‡*University of Chicago and the American Bar  
Foundation*

Trading data from the gambling market for the 2002 World Cup provide a unique window through which to test theories of market efficiency. This market provides many of the benefits of a laboratory experiment, but with much higher stakes, experienced participants, and a naturally-occurring environment. The primary drawback of the data is the relatively small number of trades. The evidence concerning market efficiency is mixed. Although markets respond strongly to goals being scored, there is some evidence that prices continue to trend higher for 10–15 minutes after a goal. We also observe systematically negative returns for bets on the pre-game favorite, consistent with the biases seen in wagering on other sports. We document the endogenous emergence of market makers. These market makers are involved in a large share of trades. Increasing from two active market makers to five or more market makers does not appear to improve the functioning of the market. On average, the market makers earn slightly negative returns, implying that other traders are able to identify situations in which market makers are setting inefficient prices.

Data limitations pose an important obstacle to testing theories of finance in real-world financial markets. Researchers typically cannot directly observe the arrival of new information to the market. The identities, characteristics, and information sets of particular traders are likewise usually unknown to the econometrician. In the face of these intrinsic data weaknesses, testing some of the most basic predictions, e.g., that efficient markets quickly and completely incorporate new information into prices, become difficult to test.<sup>1</sup>

One approach to testing whether new information is incorporated quickly and fully into markets is to analyze special cases in which new information becomes available to market participants at an identifiable point in time, e.g. announcements about the money supply (Pearce and Roley 1985, Chen *et al.* 1986), earnings (Beaver 1968), layoffs (Worrell *et al.* 1991), etc. One limitation of such tests is that some actors may have access to the information (or some portion of it) in these announcements ahead of time, either through research or insider information. Jarrell and Poulsen (1989), for instance, document large price run-ups for tender offer targets in advance of the offers being made.

A second approach researchers have used to circumvent the complexities of real-world financial markets has been to create artificial markets in controlled laboratory settings. Chamberlain (1948) and Smith (1962),

---

\*Corresponding author. Ricard Gil, Assistant Professor, Economics Department, 401 E2 Building, University of California, Santa Cruz, CA 95064. Email: rgil@ucsc.edu; Steven Levitt, Department of Economics, University of Chicago, 1126 E. 59<sup>th</sup> Street, Chicago, IL 60637. Email: slevitt@uchicago.edu

represent early applications of this approach, finding that artificial markets did not reach the equilibrium prices predicted by theory, or did so only under restrictive conditions. More recently, Plott and Sunder (1988) and List (2004) find evidence more consistent with theory. One drawback of these experimental studies is the artificiality of the setting, as argued in Levitt and List (2007a, 2007b). Subjects newly exposed to an experimental market may only learn slowly how to optimally trade, particularly if these subjects have limited experience trading in other financial markets or the rules of thumb used in naturally occurring markets are not rewarded in the lab. In addition, the amount of money at risk in experiments is small. Consequently, the cost associated with sub-optimal trading behavior is likewise small.

In this paper, we follow a third approach that combines some of the most desirable features of the research strategies described above. Namely, we analyze trading in a non-traditional, but real-world, market for outcomes of 2002 World Cup soccer matches.<sup>2</sup> The data were provided by the internet trading site [www.intrade.com](http://www.intrade.com). Unlike typical betting markets in which all bets must be placed prior to the start of an event and a bookmaker takes one side of the bet, the market we analyze is a true financial market with real-time trading during the course of a match and contracts with traders on both sides of each transaction.

This market provides many of the benefits of an experimental setting. First, the primary innovation in information in this market is the scoring of a goal. Goals are relatively rare (2.32 per game on average, or about one goal for every 40 minutes of play) and have a large impact on the asset's final value. Because these games are televised, all traders have virtually simultaneous access to new information in the form of goals, and it is also observable to the econometrician. Second, we know the identities of all traders in our data and can match IDs across trades and assets. Like standard financial markets, these traders have self-selected and presumably are knowledgeable soccer fans with experience betting on game outcomes. Third, we have not just completed transactions, but also the full set of bid and ask offers, even if these offers are cancelled or expire prior to execution. Fourth, the asset we examine has a pre-determined terminal date (i.e. the end of the game), unlike equities and currencies. This eliminates the need to take into account expectations about future events such as the flow of future dividends when pricing the security. Finally, the amount of money at stake, while not large in any absolute sense (a total of \$1.5 million worth of contracts were traded), is far greater than in experimental settings.

Using these data, we are able to test a wide array of predictions about how markets function. In favor of efficient markets, we find that the prices in soccer betting markets move sharply when a goal is scored, but that prices 10–15 minutes after the goal is scored are statistically significantly higher than immediately after the goal, providing some evidence of market inefficiency. The markets do appear to be efficient in the sense that there are few pure arbitrage opportunities available by trading combinations of securities related to a particular match. On the other hand, at least in our small sample, there are systematic biases in valuations with traders overestimating the likelihood that

#### TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

the pre-game favorite will win the match. While we are guarded in our interpretation of this result given the limited number of games in the sample, it is consistent with past evidence of a bias towards favorites in sports betting markets (e.g., Golec and Tamarkin Gray and Gray 1997, Levitt 2004, Woodland and Woodland 1994).

We also document the endogenous emergence of traders acting as market makers (i.e. simultaneously offering to buy and sell a given contract at some bid-ask spread). Between two and eight market makers operate in the various games in our sample. Market makers represent at least one side of 65 percent of all trades made while the games are taking place. We find little difference in the functioning of the market (e.g. the size of the bid-ask spread) as the number of market makers increases. Market makers overall make losses, despite the fact that their offers to buy are typically 5 percentage points below their offers to sell at any point in time. This implies that the traders taking the other side of trades with market makers have some success in identifying those instances where market makers are mispricing.

The remainder of the paper is structured as follows. Section II describes the data and presents summary statistics. Section III tests various predictions of the efficient market hypothesis and theories of trading volume. Section IV analyzes market maker behavior. Section V concludes.

## DATA DESCRIPTION

The data we use are from [www.intrade.com](http://www.intrade.com), an internet website offering trading in thousands of sports, financial, and political contracts annually.<sup>3</sup> Intrade.com is not a bookmaker and does not take positions in the securities. Rather, they create securities based on publicly observable outcomes (e.g. which team will win a game, how many inches of snow will fall in New York in December, who will win the U.S. presidential election). Traders with accounts at the web site then post bid and ask offers, with intrade.com serving as the clearinghouse for matching buyers to sellers. In return for their services, intrade.com charged a transaction cost of eight cents per trader for each \$10 contract traded.<sup>4</sup>

The particular contracts we focus on are futures contracts on the outcomes of soccer matches played as part of the 2002 World Cup. For each match, there are three contracts offered: one that pays off if a particular team wins, another that pays off if the other team wins, and a third that pays off if it is a draw.<sup>5</sup> The outcomes of the contracts are mutually exclusive and exhaustive, so exactly one contract pays off on any game.

All trades are made online using a pre-established trading account. For each contract, market participants are able to observe in real-time a market order book that lists the full set of prices and quantities of current bids and asks for each security. The trader also observes his own positions and account balances.

For each contract offered our data set includes the full order book second by second, along with the precise time, price, volume traded, and identities

of each trader. Trading occurs somewhat infrequently in these markets – on average there are 63 trades over the life of a contract, 64 percent of which occur during the roughly two hours over which the game is being played (there are two 45 minute halves, a 15 minute intermission, a few minutes of referee determined extra time). Consequently, in much of our analysis, we aggregate trading data from continuous time up to a unit of analysis of one minute intervals. When more than one transaction occurs in a given minute on a particular security, we take quantity-weighted averages of the prices to construct our price. If no trade occurs in a particular minute, then that minute is excluded from any analysis involving prices.

We match the price data with published information on the time at which each goal is scored over the course of the game. Our trading data are recorded in Greenwich Mean Time, whereas the timing of goals is based on the minutes of play that have elapsed in the game. This introduces some possible noise into the matching between prices and the timing of goals. To minimize errors, we had two research assistants independently undertake the task of performing the matching. In almost every game, the research assistants produced identical mappings. In the cases where there was disagreement, the discrepancies were only 1-2 minutes, and the authors used their best judgment to arbitrate the appropriate choice.

Table 1 presents basic information on each of the 50 games that comprise the sample. We show the teams playing, date, final score, and whether the outcome of the game was an upset based on the pre-game prices. We also report the dollar value of contracts traded, the number of trades, the number of market makers active, the number of different individuals making at least one trade on the game, and the fraction of the dollar volume which is traded while the game is in the running (as opposed to in advance of the game). The games took place over the course of roughly a month in June 2002. The total dollar value of trades ranged from a low of \$2,030 to over \$70,000. The typical game has between one hundred and two hundred trades. There are at least two market makers in each game, with as many as eight sometimes operating. The median number of individuals trading in a particular game is 75.

## **TESTING THE EFFICIENT MARKET HYPOTHESIS**

The efficient market hypothesis makes a number of predictions for these data: (1) newly arriving information should be quickly and completely incorporated into market prices, (2) no pure arbitrage opportunities across securities should exist, and (3) expected profits should be equalized across characteristics of positions that are readily observable to market participants, such as whether a team is a favorite or an underdog. We explore these three predictions in turn.

*The efficiency with which new information is incorporated into market prices*

The hallmark of efficient markets is the ability to quickly and completely incorporate new information into asset prices. Once a soccer match is underway, the primary innovation in information is the scoring of a goal.<sup>6</sup>

TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

TABLE 1  
SUMMARY STATISTICS OF GAMES PLAYED IN THE 2002 WORLD CUP

Game	Date	Score	Upset	Volume Traded (\$)	Trades	Market Makers	Traders	% Volume Traded During Game
Portugal-USA	6/05/2002	2-3	1	47010	424	7	140	0.69
Russia-Tunisia	6/05/2002	2-0	0	18810	122	4	64	0.55
Germany-Ireland	6/05/2002	1-1	0	14590	53	2	25	0.62
Cameroon-Saudi Arabia	6/06/2002	1-0	0	30000	186	4	64	0.75
Denmark-Senegal	6/06/2002	1-1	0	14690	60	2	43	0.09
France-Uruguay	6/06/2002	0-0	0	19410	170	3	83	0.55
Sweden-Nigeria	6/07/2002	2-1	0	37530	112	2	38	0.48
Spain-Paraguay	6/07/2002	3-1	0	29800	182	2	76	0.72
Argentina-England	6/07/2002	0-1	1	12240	81	2	42	0.62
Slovenia-South Africa	6/08/2002	0-1	1	20340	121	3	48	0.61
Italy-Croatia	6/08/2002	1-2	1	27940	208	4	88	0.57
Brazil-China	6/08/2002	4-0	0	40810	246	6	103	0.36
Mexico-Ecuador	6/09/2002	2-1	0	42430	344	7	140	0.5
Russia-Japan	6/09/2002	0-1	1	24420	186	4	73	0.64
Turkey-Costa Rica	6/09/2002	1-1	0	40850	305	3	124	0.65
Portugal-Poland	6/10/2002	4-0	0	24700	144	2	63	0.61
South Korea-USA	6/10/2002	1-1	0	40900	248	3	89	0.82
Belgium-Tunisia	6/10/2002	1-1	0	71520	445	8	136	0.54
France-Denmark	6/11/2002	0-2	1	34540	283	6	111	0.68
Ireland-Saudi Arabia	6/11/2002	3-0	0	28010	151	7	85	0.35
Germany-Cameroon	6/11/2002	2-0	0	29630	149	4	72	0.11
Senegal-Uruguay	6/11/2002	3-3	0	8590	31	4	20	0.42
Paraguay-Slovenia	6/12/2002	3-1	0	12790	53	2	38	0.13
Argentina-Sweden	6/12/2002	1-1	0	5640	49	3	40	0.1
England-Nigeria	6/12/2002	0-0	0	32560	204	4	84	0.59
Spain-South Africa	6/12/2002	3-2	0	59900	358	7	115	0.46
Croatia-Ecuador	6/13/2002	0-1	1	21820	107	4	70	0.17

TABLE 1  
CONTINUED

Game	Date	Score	Upset	Volume Traded (\$)	Trades	Market Makers	Traders	% Volume Traded During Game
Italy-Mexico	6/13/2002	1-1	0	43850	370	4	123	0.67
Turkey-China	6/13/2002	3-0	0	9030	64	4	39	0.37
Brazil-Costa Rica	6/13/2002	5-2	0	13720	89	4	49	0.62
USA-Poland	6/14/2002	1-3	1	6240	63	3	54	0.09
Portugal-South Korea	6/14/2002	0-1	1	64230	473	8	135	0.62
Russia-Belgium	6/14/2002	2-3	1	2030	21	2	17	0.01
Japan-Tunisia	6/14/2002	2-0	0	23350	130	4	74	0.35
England-Denmark	6/15/2002	3-0	0	31460	169	8	87	0.17
Germany-Paraguay	6/15/2002	1-0	0	10410	104	4	57	0.82
Spain-Ireland	6/16/2002	1-1 (3-2)	0	29930	190	6	92	0.29
Sweden-Senegal	6/16/2002	1-2 (GG)	1	30090	175	4	58	0.61
Brazil-Belgium	6/17/2002	2-0	0	29650	304	5	134	0.4
Mexico-USA	6/17/2002	0-2	1	33120	161	4	79	0.61
Japan-Turkey	6/18/2002	0-1	1	23760	175	5	80	0.45
Italy-South Korea	6/18/2002	1-2 (GG)	1	39230	370	8	139	0.39
Brazil-England	6/21/2002	2-1	0	52820	295	8	99	0.31
Germany-USA	6/21/2002	1-0	0	50190	328	8	141	0.27
Senegal-Turkey	6/22/2002	0-1 (GG)	1	16460	110	2	46	0.62
Spain-South Korea	6/22/2002	0-0 (3-5)	1	17420	104	2	44	0.68
Germany-South Korea	6/25/2002	1-0	0	49820	277	5	112	0.41
Brazil-Turkey	6/26/2002	1-0	0	45300	221	9	115	0.24
South Korea-Turkey	6/29/2002	2-3	1	22890	150	7	70	0.31
Brazil-Germany	6/30/2002	2-0	0	10210	71	2	42	0.63

There were a total of 64 games in the 2002 World Cup. Only 50 games are included in our sample because 14 of the games had no activity in the market. First Round games can end in a tie. Playoff games go into an overtime period if they are tied at the end of regulation play. Score reflects the number of goals each team scored. GG denotes that the game went into overtime and was won with a golden goal in an extra period. If there were no goals in the overtime period, the teams went to a shootout with penalty kicks. The results of the penalty kicks are shown in parentheses. Upset shows when the pregame underdog (as determined by the prices of the contracts immediately before the game began) beat the pregame favorite.

TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

To test the response of these markets to a goal being scored, we analyze the pattern of prices in the 30 minute window bracketing each score (15 minutes before and 15 minutes after the goal) by estimating regressions of the form:

$$(1) \quad Price_{m,c,w,t} - \overline{Price}_{m,c,w,pre} = \sum_{g=-15}^{15} \beta_g Goal_{m,c,w,t+g} + \varepsilon_{m,c,w,t}$$

where  $m$  identifies the match,  $c$  refers to the particular contract on the match (i.e. pre-game favorite to win, pre-game underdog to win),  $w$  corresponds to the particular window of 30 minutes bracketing a score, and  $t$  represents the minute of the match.<sup>7</sup> In order to anchor the prices for each contract and window around a common value, we normalize all prices by subtracting the mean price at which the contract traded five to fifteen minutes prior to the goal being scored.<sup>8</sup> The thirty  $\beta$  coefficients represent the minute-by-minute price deviations in the 15 minute time period leading up to and following a goal being scored.<sup>9</sup> *Goal* is an indicator variable corresponding to a particular minute which is equal to one when a goal is scored in the team's favor, negative one when the opposing team scores, and zero otherwise. We estimate Default (1) using weighted least squares, with the weights determined by the volume of contracts traded in that minute of the match. Parts of the game that are not within 15 minutes of a goal are excluded from the regression. If multiple goals are scored within a 30 minute period, the same observation will be included more than once since it falls into multiple goal windows. We cluster by game and contract type to take into account the resulting correlation across observations. If a goal is scored in the last 15 minutes of the game, the sample window will be truncated by the end of the game. Depending on the regression, we will sometimes limit the sample to a subset of goals, e.g. goals scored by teams with a one-goal lead immediately before the goal occurs.

The results of estimating Default (1) for the full sample of goals are shown in Figure 1. We show the  $\beta$  estimates along with 95 percent confidence intervals. Prices are flat on average in advance of a goal and rise sharply just after the goal is scored.<sup>10</sup> The magnitude of the price change accompanying a goal is approximately 20–30 on a scale that runs from 0 to 100, meaning that scoring a goal on average changes the bettors' estimates of the likelihood that this team will ultimately win by 20–30 percentage points and lowers the likelihood the opposing team will win by that amount.<sup>11</sup> The estimates are very precise before the goal is scored (as evidenced by the relatively narrow 95 percent confidence interval) and somewhat less precise after a goal is scored because of the heterogeneous impact on price of goals in different situations.

There is some visual evidence that the initial price changes are not as large as those 10–15 minutes after the goal is scored, implying that the new information is not fully and immediately incorporated. Indeed, we can reject the null hypothesis that the prices in the first few minutes after a goal is scored are equal to the later prices.<sup>12</sup> A strategy of buying contracts in the period 1–3 minutes after a goal is scored, and then liquidating these positions 13–15

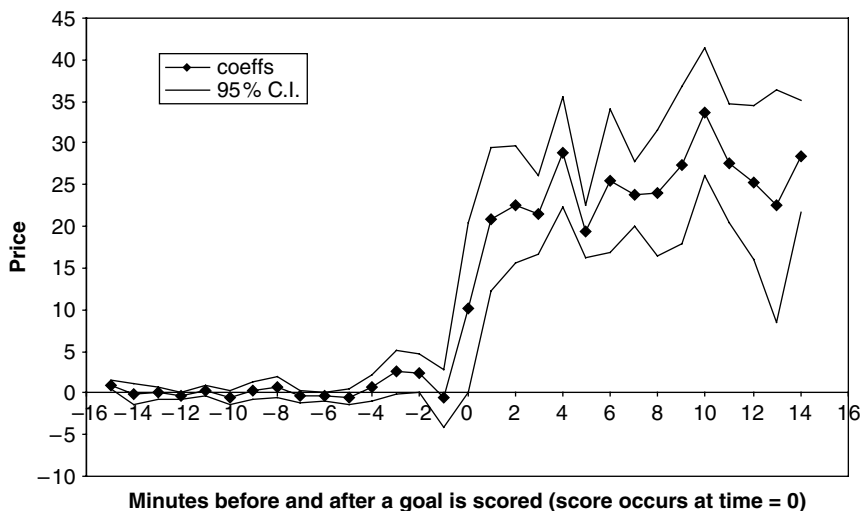


FIGURE 1. Price Response to a Goal Scored.

minutes after that same goal, would have yielded a positive return on investment of roughly 4 percent before transaction costs and 2.4 percent after transaction costs.

*The frequency with which arbitrage opportunities arise across related securities*

In efficient markets, one does not expect arbitrage opportunities to be large or persistent. In games in which a draw is possible, there are two combinations of trades that can potentially yield arbitrage profits: (1) going long all three contracts (the pre-game favorite to win, the pre-game underdog to win, and the draw), or (2) shorting all three contracts. In the first case, one gets paid off on exactly one contract with certainty. In the second instance, the arbitrageur is guaranteed to have to pay off exactly one of the three contracts. If the total of the ask prices on the three contracts (plus transaction costs) is less than 100 at any point in time, an arbitrage opportunity is available. And similarly, if the three bid prices (minus transaction costs) are more than 100, then the second arbitrage opportunity is present.

Table 2 reports the frequency, persistence, and magnitude with which such arbitrage opportunities arise. Column 1 shows that there were a total of 3 instances we identify in which an arbitrage opportunity was present. The arbitrage chances lasted, on average, for less than 15 seconds, as demonstrated in Column 2. Thus, during the playing of these 50 games, there was a total of less than 20 minutes out of roughly 100 hours of play in which an arbitrage opportunity was available, or about one-third of one percent of the time. Given that trading in these markets is relatively sparse (not a single trade occurs in about two-thirds of the game minutes we observe), it is interesting that these arbitrage opportunities so quickly correct themselves.



TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

TABLE 2  
SUMMARY OF ARBITRAGE OPPORTUNITIES

Position:	Unique arbitrage opportunities	Average length of time opportunity exists (in seconds)	Total value of arbitrage opportunities (\$)
Long on all three	17	12.53	\$100.38
Short on all three	56	14.55	\$675.98

Arbitrage opportunities take into account transaction costs associated with buying or selling a contract. The unit of observation is game x second. The unique number of arbitrage opportunities counts each arbitrage opportunity only once (an opportunity that lasts 10 seconds will only be counted once). The average length of time the arbitrage opportunity exists is conditional on this seconds having an arbitrage opportunity. The total value of the arbitrage opportunities displays the total amount of money that could have been made by fully exploiting the arbitrage opportunities.

The rapidity of the market adjustment is even more remarkable in light of the fact that the total profit that an arbitrageur would have made by exploiting all of these opportunities totals to less than \$800.

*Equal returns across readily observable security characteristics*

In an efficient market, one would not expect readily observable characteristics of securities to be systematically associated with excess returns.<sup>13</sup> Table 3 summarizes the dollar value invested, realized profits (after transaction costs), and return on investment for the various types of contracts. In each case, we report results for both investors who are long the contract and those that are short the contract. The first row includes all contract types. The total dollar value of investment at risk across all trades (both long and short) was \$1.447 million. Well more than half of the money at risk was for those who were short the position (\$875,307); because there are three contracts on each game and only one contract will pay out at the end of the game, short sellers typically have to put up more money than those going long. Investments that were long the securities lost a total of almost \$55,000. Short sellers made over \$30,000. The difference between these totals (\$23,147) reflects the transaction cost associated with Intrade's commission of 1.6 percent per contract (0.4 percent of the face value of each \$10 contract for each trader both in opening and closing a position). The profits associated with taking the short position are difficult to explain. There are no institutional obstacles associated with short selling in this context – any trader can readily take a short position. Being short does generally require a greater investment (because only one of three contracts pays off), but it is hard to believe this can account for the striking pattern in the data.

The next three rows divide the sample into the three different contract types: the pre-game favorite to win, the underdog to win, and a draw to occur. Long positions on the favorite to win yield negative returns of almost \$50,000, with corresponding profits of almost \$38,000 for those short the pre-game favorite to win (the gap between these two is again the commission charged on transactions).

TABLE 3  
REALIZED PROFIT BY CONTRACT TYPE AND POSITION

Contract:	Total Investment (\$)		Total Return (\$)		ROI(%)			
	Total	Short	Total	Short	Long	Short		
		Long	Long	Short				
All	1446680	571373	875307	23147	-54797	31650	-9.59	3.62
Favorite to win	627920	339803	288117	10047	-47897	37850	-14.10	13.14
Underdog to win	485540	127839	357701	7769	-9004	1235	-7.04	0.35
Draw to win	333220	103731	229490	5332	2104	-7435	2.03	-3.24

For each event (game), there were three contracts we included in the sample: a contract that pays if Team A wins, a contract that pays if Team B wins, and a contract that pays if Teams A & B draw. Total Investment and Total Return are dollar figures and are broken down by whether the trader offers to sell (Short) or offers to buy (Long). Total Return takes into account the fee that In Trade takes on each transaction or expired contract of fee. The fees relevant for our data amount to 8 cents per person per contract which amount to 16 cents per trade (In Trade has since lowered its fees). ROI is the return on investment, defined as the total return divided by the total investment. A team is described as the Favorite if the contract of it to win was trading at a higher price than either of the other two contracts on that game the minute before the game started.

TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

TABLE 4  
WINS, TIES, AND LOSSES FOR MATCH FAVORITES IN PREVIOUS WORLD CUPS

Year	Won	Tied	Lost
1994 (USA)	49.02 (N = 25)	23.53 (N = 12)	27.45 (N = 14)
1998 (France)	54.69 (N = 35)	25.00 (N = 16)	20.31 (N = 13)
2002 (Korea & Japan)	60.94 (N = 39)	21.88 (N = 14)	17.19 (N = 11)

This table reports the number and percentage of team s that lost, tied, or won their games conditional on being the favorite as ranked by FIFA. A team was ranked as the favorite if its FIFA ranking was higher than that of its opponent prior to the beginning of each World Cup. For this table only, we use the FIFA ranking to determine favorites in 2002 for consistency with earlier World Cups. The rankings vary over time. Note that the 1994 World Cup only had 51 games while the 1998 and the 2002 World Cups both had 64 games.

The availability of positive returns shorting favorites runs contrary to efficient market predictions. One possible explanation for the findings is that they are an aberration resulting from a small sample of games. If an unusually large number of upsets and an unexpectedly small number of draws occurred in the 2002 World Cup, such a result would arise in sample. Table 3, which displays game outcomes for the three most recent World Cups, however, calls into question this explanation for the findings. Favored teams won a greater share of games in the 2002 World Cup (our sample) than in either the 1994 or 1998 tournaments.<sup>14,15</sup>

For the other two contract types, returns are much closer to being equalized across short and long. Being short the underdog yields a small profit, as does being long the draw.

## THE ROLE OF MARKET MAKERS

The markets we are analyzing are relatively thinly traded, making liquidity an issue. One interesting feature of these markets is the endogenous emergence of market makers who provide liquidity by offering to either buy or sell a substantial volume (at least \$500 worth) of a given contract at a particular point in time. These market makers are not affiliated with Intrade and do not get special treatment vis-à-vis other traders. They are simply individual traders who elect to provide market liquidity in pursuit of profits. Market makers are on at least one side of the trade in 65 percent of the total contracts traded and they represent approximately 50 percent of the money invested.

Table 5 provides information about the distribution of numbers of market makers, trading activity, and bid-ask spreads by contract. For every contract we observe, there are at least two traders who take a market maker role. On over 20 percent of the contracts there are five or more market makers in action (eight is the maximum). The second row of the table reports the average volume traded on each contract, by number of market makers. Trade volumes are an increasing function of the number of market makers. This correlation may reflect the fact that additional market makers increase liquidity in the markets. Alternatively, it may simply be that the same factors that generate

TABLE 5

VALUE OF TRADES AND BID-ASK SPREAD INFORMATION BY NUMBER OF MARKET-MAKERS

	Number of Market Makers active in the Contract			
	2	3	4	5 +
Number of Contracts	46	37	32	35
Average \$ Trade Volume	5643	9180	11149	14019
Median Bid-Ask Spread	4	3	4	3
Mean Bid-Ask Spread	5.71	5.11	6.15	4.34

Data are by minute and contract. The Bid-Ask spread is the difference between the highest bid (buy) offer and the lowest ask (sell) offer at each minute of the game (the actual measurements are taken at exactly 1 minute into the game, exactly 2 minutes into the game, etc.). There were no contracts with fewer than 2 market makers or more than 8 market makers. The threshold for being considered a market maker is simultaneously offering to buy and sell at least 50 units of a single contract.

trading interest also encourage more market makers to enter. Analysis of the median and mean bid-ask spreads on the contracts suggest that the latter explanation is likely more important. The median bid-ask spread varies between three and four and it is not clearly declining with the addition of more market makers.<sup>16</sup> The mean bid-ask spread is highest when there are four market makers, and lowest when there are five or more market makers. In sum, it appears that the contracts with two market makers provide about as smoothly as those with more market makers.

Table 6 shows returns to traders by the total number of trades they make, with market makers broken out separately. More active traders do better than those who make fewer trades, although this may be due in part to selection: if a trader has done poorly thus far betting on the World Cup, he may be less likely to continue trading. The interesting exception to the pattern of frequent traders doing better is for market makers. Although market makers are extremely active traders, their net profits are negative, even before transaction costs.

If the bid-ask spread set by market makers is centered on the true probabilities and the spread offered by the market makers is 4 percent, then the

TABLE 6

DISTRIBUTION OF PROFITS BY TRADER ACTIVITY AND NUMBER OF TEAMS BET ON

By Number of Trades:	Total Investment (\$)	Total Return (\$)	ROI (%)
Small (0-49)	177532	-16877	-9.51
Medium (50-99)	186864	-539	-0.29
Large (> 99)	378162	11892	3.14
Market Makers	704122	-17622	-2.50

The unit of observation is a trader. Market makers are not included in the Number of Trades categories except for the Market Makers category. A person is considered to be a market maker if she has concurrent offers to buy and sell more than 50 units of a given contract. Market maker status can vary across contracts for the same person (a person can be a market maker on one contract and not a market maker on a separate contract). Total Investment and Total Return are dollar figures. Total Return takes into account the fee that In Trade takes on each transaction or expired contract offer. The fees relevant for our data amount to 8 cents per person per contract (In Trade has since lowered its fees). ROI is the return on investment, defined as the total return divided by the total investment.

#### TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

expected market maker profit would be 2 percent per trade, minus a 0.8 percent transaction fee, for a net profit rate of 1.2 percent per trade. The total dollar value of trades by market makers of about \$700,000 would imply an expected profit of \$8,400. In practice, however, market makers lose \$17,622 in the sample, for a return on investment of  $-2.5$  percent. This suggests that the bids and asks posted by the market makers are not centered around the true probabilities, and that on average the bettors on the other side of the trade have better estimates of the true probabilities than the market makers.<sup>17</sup> Market makers do particularly badly just before and after a goal is scored – about 30% of their total losses arise in the minute before and minute of a goal. This suggests that in some cases other gamblers are able to take advantage of the arrival of new information before the market makers are able to cancel their posted bid-ask spreads.

## CONCLUSION

Trading in securities tied to the outcome of World Cup matches provide a unique window onto the working of real-world markets. These contracts have many of the virtues of laboratory experiments on markets: the security traded is simple, it is easy to identify information innovations and all market participants observe these information shocks simultaneously. The World Cup contracts also have important advantages over the lab. The amount of money at stake is much greater than most laboratory settings, the markets are naturally occurring and not subject to the distortions that come with the excess scrutiny of the lab, and the participants self-select into these markets in a way that is presumably analogous to other actual markets (Levitt and List 2006). Our analysis of these markets suggests that the patterns observed provide only mixed support for the efficient market hypothesis. Although prices jump immediately after goals are scored, prices are higher still 10–15 minutes after the goal. Moreover, we observe systematically negative returns for bets on favorites, which are not predicted in an efficient market. This finding does not appear to be an artifact of the relatively small sample of matches that we observe, and mirrors similar patterns in other sports wagering markets. In favor of market efficiency, arbitrage opportunities across inter-related markets are few and far between. Finally, we observe the endogenous emergence of market makers providing liquidity in these markets. Surprisingly, these market makers generate negative profits, yet remain prevalent in the data.

In addition to the puzzles highlighted above, these markets also raise another important question worthy of economic consideration. The structure of this trading market is fundamentally different than the standard sports wagering contract in which a bookie takes one side of a bet with a bettor on the other. To compensate for the resulting risk, the bookie's commission, known as the vigorish, is typically 5 percent of the value of the bet. The commission on trades at Intrade, in contrast, are less than 2 percent of the value of the bet. Further, the Intrade contracts have the added virtue of being traded at any

point in time, even during the event, whereas the standard bet cannot be reversed once it is placed with a bookie. In spite of these advantages, a total of less than \$1.5 million in trades were made at Intrade in the 2002 World Cup, a trivially small fraction of overall betting on the event. A partial answer to that question no doubt lies in the fact that the Intrade website was relatively new at that time. Nonetheless, more recent data from other sporting events suggests that while this form of wagering is gaining increased acceptance among bettors, the share of all bets done in this format is far less than that which economists would predict given the clear advantages that these markets provide relative to traditional forms of bookmaking.

#### ACKNOWLEDGEMENTS

We would like to thank Lars Hansen and John List for extensive discussions that have greatly improved the paper, as well as participants in the May 2007 conference on The Growth of Gambling and Prediction Markets held at University of California at Riverside. Ethan Lieber, Michael Roh, and Kenneth Ward provided stellar research assistance. We are grateful to John Delaney of Intrade.com for providing the data used in this paper. Financial support for this research was provided by the National Science Foundation and the Sherman Shapiro Research Fund.

#### NOTES

1. There is, of course, an enormous literature testing various predictions of the efficient market hypothesis. An early survey of this work is Fama (1970). For two divergent reviews, see Malkiel (2003) and Shiller (2003).
2. Camerer's (1998) analysis of the impact of temporary interventions in horse racing betting market, while very different from our own study in the details, also has the feature of merging elements of experimental studies with a real-world market.
3. Leigh et al. (2003) use data from intrade.com's sister site, tradesports.com, to analyze the response of financial markets to the threat of war in Iraq. Hartzmark and Solomon (2006) analyze tradesports.com data to test efficiency in football betting markets. Wolfers and Zitzewitz (2004) summarizes the existing literature on prediction markets.
4. At the time of our sample, each party to a transaction pays four cents per contract to open a position and four cents per contract to close a position. Positions that expire worthless are also charged a four cent closure fee. The fee structure has since changed at intrade.com, and the commissions charged are now lower. Note that these transaction costs are much smaller than the typical "vig" at traditional sports books, where bettors must lay 11 units to win 10 on gambles that are 50–50 bets.
5. Because the contracts are for regulation time only (including injury time), in later rounds where the game cannot end in a draw, we still observe a contract that pays off for a draw.
6. Of course, the passage of time without a goal scoring also reveals information. Since the expected number of goals scored per minute is only 0.026, the scoring of a goal is a much greater shock to the outcome of a game than failing to score a goal at any particular time.
7. We label the minute the goal is scored as minute number one; we do not have a minute zero in our analysis, instead jumping directly from minute negative one to minute one.
8. This is similar to including match-window fixed effects, except that it only conditions on the pre-goal prices. To the extent that the price impact of a goal varies by game situation, normalizing in this way provides an easier means of interpreting the impact of a goal.
9. Note that, regardless of the process that generates goals (e.g. perhaps the opponent is more likely to score a goal shortly after the first team scores, or one team is far more likely than the other team to score throughout the game), efficient markets predicts that all of the post-goal  $\beta$ 's will be equal. A number of previous readers have been confused by this point.
10. Note that minute zero on the graph is the minute the goal is scored, so some trades occur before the score and others occur after.
11. In the specification in Figure 1, we have constrained the coefficient on price to have the same magnitude (but the opposite sign) when a team scores a goal versus is scored upon. Because of the

## TESTING THE EFFICIENCY OF MARKETS IN THE 2002 WORLD CUP

possibility of a draw, these two values need not be identical. In later figures, we report results separately for goals in favor and goals against.

12. Although not shown graphically, we have also analyzed the price response to a goal being scored for a range of subsets of the data. The impact on the price of a team to win the game is similar when the team scores a goal to achieve a one goal lead or a two goal lead (on average the price change is 30 percentage points). When the team scores to tie the game, the price impact is smaller (slightly below 20 percentage points). The impact of a goal is not particularly sensitive to the price of the contract before the goal is scored. For almost all of these sub-groups we can reject the null hypothesis that the price changes are immediately and completely incorporated after a goal is scored.
13. In financial markets, if observable characteristics are correlated with risk factors, then expected returns may vary, but no systematic correlation between observable factors and risk-adjusted returns should exist.
14. In the earlier tournaments, we do not have prices from the betting markets, so we define the favorite to be the team with the higher FIFA ranking prior to the beginning of World Cup play to be the favorite. In the 2002 World Cup, this algorithm would correctly pick the team that bettors call the favorite in 44 of the 50 games in our sample. For consistency, we use FIFA rankings to define favorites and underdogs in 2002 in Table 4, although this is not our definition elsewhere in the paper.
15. The results concerning bettor bias towards favorites, although far from definitive given our small sample of games, is consistent with past evidence of such biases in wagering on other sports (Levitt 2004).
16. Although not shown in the table, when contracts are broken down separately into 5, 6, 7, or 8 market makers, the median bid-ask spread is 3 for each of those cases.
17. Further evidence that the market makers' posted prices are not centered around the true probabilities comes from the fact that about 6.6% of trades take place between two market makers, when one's bid price is above the other's ask price.

## REFERENCES

- William Beaver 'The information content of annual earnings announcements' *Journal of Accounting Research* (1968) 6 67–92.
- Colin Camerer 'Can asset markets be manipulated? A field experiment with racetrack betting' *Journal of Political Economy* (1998) 106(3) 457–482.
- Edward Chamberlain 'An experimental imperfect market' *Journal of Political Economy* (1948) 56(2) 95–108.
- N. Chen, Roll Richard and Ross Stephen 'Economic forces and the stock market' *Journal of Business* (1986) 59 383–403.
- Eugene Fama 'Efficient capital markets: A review of theory and empirical work' *Journal of Finance* 25(2) 383–417.
- J. Golec and M. Tamarkin 'The degree of price inefficiency in the football betting markets' *Journal of Financial Economics* (1991) 30 311–325.
- P. Gray and S. Gray 'Testing market efficiency: Evidence from the NFL sports betting market' *Journal of Finance* (1997) 52 1725–1737.
- Sam Hartzmark and Solomon David *Are Gamblers Loss Averse? Efficiency and the Disposition Effect in N.F.L. Prediction Markets*. unpublished manuscript, University of Chicago (2006).
- Gregg Jarrell and Poulson Annette 'Stock trading before the announcement of tender offers: Insider trading or market anticipation?' *Journal of Law, Economics, and Organization* (1989) 5(2) 225–248.
- Andrew Leigh, Wolfers Justin and Zitzewitz Eric *What do Financial Markets Think of War in Iraq?* NBER Working Paper No.9587 (2003).
- Steven Levitt 'How do markets function? An empirical analysis of gambling on the National Football League' *Economic Journal* (2004) 114(495) 2043–2066.
- Steven Levitt and List John 'On the generalizability of lab behaviour to the field' *Canadian Journal of Economics* (2007a) 40(2) 347–370.
- Steven Levitt and List John 'What do laboratory experiments measuring social preferences tell us about the real world?' *Journal of Economic Perspectives* (2007b) 21(2) 153–174.
- John List 'Testing neoclassical competitive theory in multilateral decentralized markets' *Journal of Political Economy* (2004) 112(5) 1131–1156.
- Burton Malkiel 'The efficient market hypothesis and its critics' *Journal of Economic Perspectives* (2003) 17(1) 59–82.

- Douglas Pearce and V. Vance Roley 'Stock prices and economic news' *Journal of Business* (1985) 58(1) 49–67.
- Charles Plott and Sunder Shyam 'Rational expectations and the aggregation of diverse information in laboratory security markets' *Econometrica* (1988) 56(5) 1085–1118.
- Robert Shiller 'From efficient markets theory to behavioral finance' *Journal of Economic Perspectives* (2003) 17(1) 83–104.
- Vernon Smith 'An experimental study of competitive market behavior' *Journal of Political Economy* (1962) 70(2) 111–137.
- Justin Wolfers and Zitzewitz Eric 'Prediction markets' *Journal of Economic Perspectives* (2004) 18(2) 107–126.
- Linda Woodland and Woodland Bill 'Market efficiency and the favorite-longshot bias: The baseball betting market' *Journal of Finance* (1994) 49 269–279.
- D. Worrell, W. Davidson, III and V. Sharma 'Layoff announcements and stockholder wealth' *Academy of Management Journal* (1991) 34 662–678.